



*Super*TRAK⁶⁶_{TM}

User's Manual

Version 1.0

Copyright

Copyright by Promise Technology, Inc. ("Promise"), 2000. No part of this manual may be reproduced or transmitted in any form without the expressed, written permission of Promise.

Trademarks

All trademarks are the property of their respective holders.

IMPORTANT DATA PROTECTION INFO

You should back up all data before installing SuperTrak66. Promise is not responsible for any loss of data resulting from the use, disuse or misuse of this or any other Promise product.

Notice

Although Promise has attempted to ensure the accuracy of the content of this manual, it is possible that this document may contain technical inaccuracies, typographical, or other errors. Promise Technology, Inc. assumes no liability for any error in this publication, and for damages, whether direct, indirect, incidental, consequential or otherwise, that may result from such error, including, but not limited to loss of data or profits.

Promise provides this publication "as is" without warranty of any kind, either express or implied, including, but not limited to implied warranties of merchantability or fitness for a particular purpose.

The published information in the manual is subject to change without notice. Promise reserves the right to make changes in the product design, layout, and driver revisions without notification to its users.

Table of Contents

Introduction	1
1.1 Promise Technology, Inc.	1
1.2 What is SuperTrak66?	1
1.3 Key Features	2
1.4 System Requirements	3
1.5 Operating System Support	3
Getting Started	5
2.1 Unpacking SuperTrak66	5
SuperTrak66 Controller Card	6
Cables	6
2.2 Quick Installation Checklist	7
Hardware Installation	7
Software Installation	7
2.3 Hardware Installation	8
2.4 Software Installation	9
BIOS Utility	9
Network Connections	9
Driver Installation for Existing Windows NT 4.0 System	11
Driver Installation During New Install of Windows NT 4.0	12
SuperCheck Installation	13
SuperBuild™	15
3.1 Setup Task Quick Reference	15
Setup Checklist	15
Common Tasks	15
Entering the BIOS SuperBuild utility	16
Creating a disk array from the Auto Setup screen	16
Creating an array from the View / Define Array screen	16
Deleting an array from the Delete Array screen	16
Setting the Halt On Error option	16
3.2 SuperTrak BIOS Reference	17
Main Menu	18
3.3 Auto Setup	19
Auto Setup Options Menu	19
Auto Setup Configuration	20
3.4 View Drive Assignments	21

3.5 View / Define Array	22
Define Array Definition Menu	23
Creating a RAID 0 Array	24
Creating a RAID 1 Array	24
Creating a RAID 3 Array	26
Creating a RAID 5 Array	27
Creating a Spanning Array	27
"Hot" Spare Drives	28
Setting a Disk Array As Bootable	28
Disk Array Recognition Order	29
Saved Disk Array Information	29
3.6 Delete Array	29
3.7 Controller Configuration	30
SuperCheck™Setup	32
4.1 SuperCheck Installation	34
Installation Scenarios	34
Common Component Installation	36
4.2 SuperCheck Administration	41
Creating Message Servers & Agents	41
Adding/Deleting New Users	41
Setting Up A New Array	41
Array Administration	42
Setting Up Email Notification	42
Identifying Problems	43
4.3 SuperCheck Quick Reference	44
To Create A New Message Server	44
Creating A New User	44
Deleting an Existing User	45
Accessing Message Agent	45
Creating a New Array	45
Deleting an Existing Array	46
Setting Up E-mail Alert Notification	46
Adding an User to an Email Receiver List	46
Removing an User From an Email Receiver List	47
Adding/Removing Events to/from Email Alert Notification	47
Using SuperCheck™	48
5.1 Main Window	48
Using Tree View	49
Using Object View	51
Using Information View	51
Using Status Bar	52

5.2 SuperCheck Pull-Down Menu/Toolbar	52
Using Pull-Down Menu Items	52
Using View Pull-Down Menu	53
Using Connection Pull-Down Menu	53
Using Preference Pull-Down Menu	54
Using Help Pull-Down Menu	54
Using Help Topics	55
Using SuperCheck Toolbar	57
5.3 Console Functions	59
Creating A New Message Server	59
Viewing Console Object icons	60
5.4 Message Server Functions	61
Deleting A Message Server	61
Disconnecting A Message Server	61
Connecting The Message Server	62
Switching to Another Server	63
Viewing Message Server Object icons	64
Viewing Message Server Information	64
5.5 Message Agent Functions	66
Accessing Message Agent	66
Creating A New User	66
Setting Up E-mail Alert Notification	68
Adding an User to an Email Receiver List	68
Removing an User From an Email Receiver List	69
Adding/Removing Events to/from Email Alert Notification	69
Message Agent Information View	71
5.6 User Management Functions	75
Creating A New User	75
Deleting an Existing User	75
Changing Passwords	76
Assigning User Rights	76
Viewing User/Admin Object icons	76
User Information View	77
5.7 I₂O RAID Functions	79
Viewing I ₂ O RAID Object icons	79
Information View	79
5.8 IOPx Functions	80
Creating a New Array	80
Reading/Clearing Events from SuperTrak Memory Buffer	81
Using The Event Viewer	82
Changing Flush Memory Timing	82
Changing Dirty Threshold Flush Start Setting	83
Changing Dirty Threshold Flush Stop Setting	83

Viewing IOPx Object icons	83
IOPx Information View	84
5.9 Channel (chx) Functions	86
Viewing Channel Object icons	86
5.10 Hard Drive Functions	87
Hard Drive Information View	87
5.11 Enclosure Functions	88
Enabling Rebuild/Synchronization Beeper	88
Viewing Enclosure Object icons	88
Enclosure Information View	89
5.12 Array Functions	90
Deleting an Existing Array	90
Rebuilding An Array	91
Using Rebuild Wizard	92
To Stop Rebuild	94
Array Synchronization	94
Halting Synchronization	95
Turning On Read Cache	96
Turning Off Caching	96
Turning On/Off Write Back Cache	96
Viewing Cache Memory Statistics	97
Information View	98
RAID Setup	102
6.1 Application Scenarios	103
High Performance Desktop PCs Working With Large Files	103
Corporate Desktop PC Requiring Data Protection	103
Windows NT File Server	103
Mid-Sized Windows NT Application Server	104
Largest Storage Capacity Required	105
6.2 RAID Implementation	105
Striping (RAID 0)	105
Mirroring (RAID 1)	106
Striping/Mirroring (RAID 0+1)	106
Data striping with dedicated parity drive (RAID 3)	106
Block and parity striping (RAID 5)	107
Spanning	107
6.3 RAID Management & Operation	108
Critical & Offline Arrays	108
Rebuilding/Synchronizing Fault Tolerant Arrays	108
Partitioning & Formatting Arrays	108
6.4 Configuration Ideas	109

Fault Tolerance	109
Hot Spare Drive(s)	109
Troubleshooting	110
7.1 Buzzer Alarms	110
7.2 LED Display Codes	111
7.3 SuperTrak66 BIOS Error Messages	112
7.4 Installation & Runtime Problems	113
APPENDIX A	Technology Background
APPENDIX B	SuperTrak Technical Specifications
APPENDIX C	Frequently Asked Questions
APPENDIX D	Contacting Technical Support
APPENDIX E	Limited Warranty
APPENDIX F	Returning Product for Repair
APPENDIX G	Glossary of Terms
APPENDIX H	Index

Table of Figures

Figure 1: SuperTrak66 controller card.....	6
Figure 2: Typical controller card installation	8
Figure 3: Driver Installed Screen in Existing NT	11
Figure 4: Driver Installed Screen in New NT.....	12
Figure 5: SuperTrak66 BIOS startup.....	17
Figure 6: SuperBuild Array Utility Main Menu.....	18
Figure 7: Auto setup screen.....	19
Figure 8: View Drive Assignments screen.....	21
Figure 9: Define Array Menu screen.....	22
Figure 10: Define Array Definition Menu screen	23
Figure 11: Two-Drive Mirroring Dialogue Window.....	25
Figure 12: Source Disk Window for Mirroring Existing Data.....	25
Figure 13: Assigning Bootable Array.....	29
Figure 14: Delete Array Menu	30
Figure 15: Delete Array Dialogue Box.....	30
Figure 16: Controller configuration screen.....	31
Figure 17: Typical SuperCheck Component Installation	32
Figure 18: SuperCheck component installation menu	36
Figure 19: "Choose Destination Location" window.....	37
Figure 20: "Select Program Folder" window	38
Figure 21: "Start Copying Files" window	39
Figure 22: "Setup Needs the Next Disk" pop-up window	40
Figure 23: SuperCheck utility main window	48
Figure 24: Using Tree View	49
Figure 25: Using Object View	51
Figure 26: Using Information View.....	51
Figure 27: Using Status Bar	52

Figure 28: Array pull-down menu.....	52
Figure 29: Using View Pull-Down Menu.....	53
Figure 30: Using Connection Pull-Down Menu	53
Figure 31: Using Preference Pull-Down Menu	54
Figure 32: Using Help Pull-Down Menu.....	54
Figure 33: Using Help Contents Window	55
Figure 34: Using Help Find Window.....	56
Figure 35: Toolbar icons	57
Figure 36: Creating New Message Server pull-downs	59
Figure 37: Labeling New Message Server	60
Figure 38: Console Object View	60
Figure 39: Deleting Message Server pull-down	61
Figure 40: Disconnecting Message Server pull-down	62
Figure 41: Connecting Message Server pull-down	63
Figure 42: Message Server Information View	64
Figure 43: Setting Password screen	66
Figure 44: Creating New User screen.....	67
Figure 45: Setting E-mail Alert box	68
Figure 46: Setting Email ID window	68
Figure 47: Setting Email Address window.....	69
Figure 48: Removing User from E-mail alert.....	69
Figure 49: Adding/Removing Events for Email Alert.....	70
Figure 50: Scheduling Array Synchronization.....	70
Figure 51: Message Agent Console icons	71
Figure 52: Message Agent Information View.....	72
Figure 53: Alert Events & Schedule Synchronization screen	74
Figure 54: Creating New User pull-down	75
Figure 55: Deleting User menu.....	76
Figure 56: User icons in Object View.....	77
Figure 57: User Information View.....	78
Figure 58: I ₂ O RAID Object View	79
Figure 59: I ₂ O RAID Information View	79
Figure 60: Creating New Array pull-down	80
Figure 61: Labeling New Array window.....	80
Figure 62: Read/Clear Events pull-down	81
Figure 63: Events Viewer screen.....	82
Figure 64: IOPx Object View icons	83
Figure 65: IOPx Information View	84
Figure 66: Channelx's Object View	86
Figure 67: Hard drive Information View	87
Figure 68: Enclosure Object View	88
Figure 69: Enclosure Information View	89
Figure 70: Deleting Array pop-up menu	90
Figure 71: Array Pull-Down menu.....	91
Figure 72: Rebuild Wizard Step 1	92
Figure 73: Rebuild Wizard Step 2	93
Figure 74: Array Rebuild Progress Bar.....	93
Figure 75: Stop Rebuild pull-down.....	94
Figure 76: Array Synchronization pull-down.....	95
Figure 77: Halt Synchronization pull-down	95

Figure 78: Array Object View.....	97
Figure 79: Array Information View.....	98
Figure A1: RAID 0 striping interleaves data across multiple drives.....	A-2
Figure A2: RAID 1 mirrors identical data to two drives	A-3
Figure A3: RAID 0+1 striping and mirroring of two drive pairs.....	A-4
Figure A4: RAID 3 multiple drives stripe data w/ one dedicated parity drive	A-6
Figure A5: RAID 5 stripes all drives with data and parity info	A-7
Figure A6: Spanning uses full capacity of drives.....	A-8

Manual Conventions

Common identifiers:

- ↳ - Press the Enter key
- [Key] - Press the key(s) shown within the brackets
- Note:** - *Supplementary note containing important information*

Common expressions:

- "Left-click" - move the mouse cursor over the specified target, then click once with the left mouse button.
- "Right-click" - move the mouse cursor over the specified target, then click once with the right mouse button.
- "Double-click" - move the mouse cursor over the specified target, then click twice rapidly with the left mouse button.
- "Drag and drop" - position the mouse cursor over the indicated icon, depress and hold the left mouse button while moving the cursor. Release the left mouse button to drop the icon at the new location.

NOTES

Introduction

1.1 Promise Technology, Inc.

Promise Technology, Inc. was founded in San Jose in 1989 and established a proven track record for leading edge storage controller products. With an innovative product line, the company has pioneered the ATA RAID storage concept, allowing users to configure RAID arrays using low-cost Ultra ATA and EIDE drives. Promise Technology was the first to introduce caching controllers to MFM, and then IDE drives, first to fully support Ultra ATA/33 and Ultra ATA/66 drives, and the first to develop a RAID 0,1 card for Ultra ATA drives.

1.2 What is SuperTrak66?

SuperTrak66 is a hardware-based RAID controller and a Windows NT 4.0 software management suite that provides an enterprise-wide ATA RAID solution. The SuperTrak66 controller card uses the I/O architecture and employs Intel's i960RD RISC processor. When used with its companion software management suite, SuperTrak66 offers a feature-rich, secure, versatile, and expandable RAID environment that allows users and administrators to configure, manage, and monitor everything from single arrays on local systems to array networks residing in offsite locations.

At its core, SuperTrak66 provides advanced RAID management functions: creating arrays, monitoring them, keeping them online and operating at optimum efficiency. SuperTrak66 can also perform many other tasks, such as:

- create arrays at various RAID levels (0, 1, 0+1, 3, 5), depending on the application being used;
- set up a network of SuperTrak66 RAID servers (all running under different RAID levels) and monitor those servers from any workstation on the network;
- create a series of SuperTrak66 RAID networks at any number of offsite locations;
- monitor and/or repair SuperTrak66 RAID arrays using the Internet from an offsite location -- all without compromising the integrity of secure servers.

The SuperTrak66 controller card features Intel's i960RD microprocessor and works in tandem with I/O, a new firmware/software layer that moves much of the I/O workload off the main CPU and transfers it to the controller card's processor, boosting the performance of the entire system. Each SuperTrak66 card can support up to four (4) Ultra ATA/66 hard drives configured as RAID levels (0, 1, 0+1, 3, or 5). Attached drives can be "hot swapped," when used with an optional FastSwap66 or SuperSwap66 drive housings.

On the software side, SuperCheck comprises three “modules”: Console, Message Server, and Message Agent. These modules use TCP/IP connections to communicate with one another, allowing communication across a network. This allows system administrators to monitor and rebuild the SuperTrak RAID system from a local console or a remote workstation over the Internet. With controlled user group administration, secure data encryption, and flexible control options, SuperTrak66’s SuperCheck software package is an all-in-one modular solution.

1.3 Key Features

The following are key features of the Promise SuperTrak66 RAID system:

Drive Types

- Supports standard IDE drives (Ultra ATA, DMA, and EIDE)
- Supports Ultra ATA/66, Ultra ATA/33, EIDE, and Fast ATA-2 (UDMA 4/3/2/1/0, PIO 4/3/2/1/0, and DMA 2/1/0, respectively)

RAID Technology

- Data striping (RAID 0), mirroring (RAID 1), striping & mirroring (RAID 0+1), parity (RAID 3), and parity striping (RAID 5) at hardware level
- Supports automatic background rebuilds for mirrored arrays

Cache

- Complete read/write cache management
- Onboard controller card memory enhances cache size
- Supports 8MB minimum, 128MB maximum, 72-pin EDO RAM memory
- Implements read and write back caching

Hot Swap Support

- Hot swappable hard drives, allowing for easy repair/replacement

Automatic Functions

- Drive fault detection
- Transparent rebuild of fault tolerant arrays in the event of drive fault

Graphical User Interface

- Remote monitoring of all arrays from network stations
- Local maintenance of arrays with drag-and-drop support

I₂O

- Implements the I₂O architecture
- Onboard Intel i960RD microprocessor offloads I/O burden from the main system CPU, boosting performance

Optional Enclosure Management (for FastSwap66 or SuperSwap66)

- Monitors enclosure temperature, fan, and power (SuperSwap66 only)

1.4 System Requirements

- 1 bus mastering PCI slot (PCI v.2.1 compliant)
- 8MB EDO 72-pin SIMM for cache RAM (units may ship with 16MB)
- One to four Ultra ATA/66, Ultra ATA/33 or EIDE hard drive(s) for dedicated use in the RAID array
- Windows NT 4.0 (or later) on server stations; Windows NT 4.0 or Windows '95 (or later) on remote monitoring stations

1.5 Operating System Support

While SuperTrak66 was designed to specifically support Windows NT 4.x and later, users may employ other operating systems to access many of its operational features. What follows is a brief description of the different support levels available for each operating system:

Windows NT

NT 4.0 is recommended to take full advantage of all the features of the SuperTrak66 package. I/O drivers are supplied Promise Technology. The SuperTrak66 Message Server and Message Agent are NT "Services," and will only function on an NT system. Without the Message Server and Message Agents installed properly on NT stations, the SuperCheck utility will not permit maintenance and configuration of your RAID systems.

Windows 95/98

Windows 95 and its successors employ the WIN32 interface required by the SuperCheck monitoring utility software. This means that SuperCheck may be implemented on a Windows 95/98 station or remote terminal to connect to Message Servers that exist on a network.

MS DOS

Using the *SuperTrak66* int 13h BIOS interface, it is possible to access the drive data on an array to partition, format, copy files, etc., using DOS 7 or later. DOS versions earlier than 7.0 have an 8.4GB drive size limit which cannot be changed.

Other Operating Systems

Other operating systems may or may not provide an OSM driver for I/O sufficient to attach to *SuperTrak66* arrays. However, Promise is not officially supporting any alternative operating system at this time.

Getting Started

This chapter is designed to help you prepare SuperTrak66 for installation into Windows NT 4.0 system servers. Please read through this chapter carefully before attempting to install *SuperTrak66*. Users should record their current CMOS (system setup) settings before making any changes. This preventative measure is aimed at protecting such information from loss, which may happen unpredictably and/or periodically during any hardware installation.

2.1 Unpacking SuperTrak66

The SuperTrak66 package should contain the following items:

- *SuperTrak66 controller card*
- *SuperTrak66 User's Manual*
- *SuperTrak66 OS drivers & utility diskettes (3)*
- *Internal 18" Ultra ATA/66 80-wire/40-pin hard drive cables (4)*

NOTE: Unit may ship with 16MB 72-pin EDO SIMM

If any of the items are missing or appear damaged, please contact your dealer or distributor immediately.

NOTE: The *SuperTrak66 controller card* ships with two removable protective labels that cover the audible alarm buzzer device (see Figure 1) and the battery (see Figure 1) which maintains the non-volatile RAM memory. These labels must be removed prior to installing the card.

SuperTrak66 Controller Card

Seen in Figure 1, the SuperTrak66 controller card has several physical features of interest for purposes of installation: A SIMM socket for EDO memory (minimum 8MB required), 4 IDE channel connections (1 "master" drive per channel), an alarm buzzer, a battery, and 4 LEDs on the card's backplane.

NOTE: The SuperTrak66 controller card, like every other valuable part of your PC, is sensitive to static electricity. Be sure that you are properly grounded (Promise recommends that you wear an anti-static strap, or place a free hand on a grounded object), and that your PC is unplugged before installing any of the SuperTrak66 components.

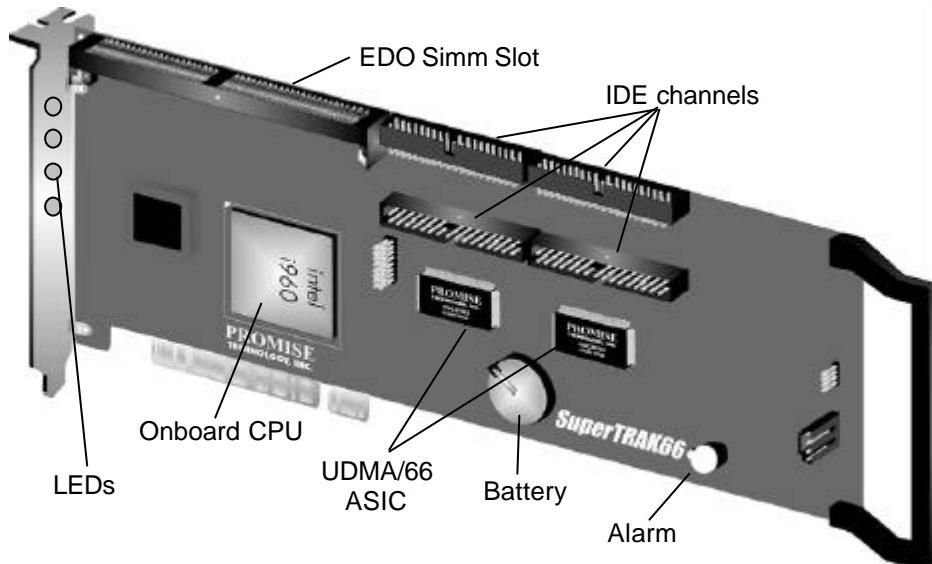


Figure 1: SuperTrak66 controller card

Cables

Ultra ATA/66 drive cables are provided with the SuperTrak66 controller card. These cables must be used with Ultra ATA/66 drives in order to realize full Ultra ATA/66 specs. They differ from older 40-wire/40-pin IDE cables by adding an extra ground wire between each data wire, resulting in an 80-wire/40-pin cable.

2.2 Quick Installation Checklist

Refer to this checklist to ensure that all hardware and software components necessary for your SuperTrak66 to operate efficiently are installed.

Hardware Installation

- At least one drive is cabled and attached to the controller.
- SuperTrak66 controller card is fully inserted into a bus mastering PCI slot and mounted with a bracket screw
- Minimum of 8MB EDO RAM is inserted into the SIMM slot on the controller card
- A 3V (CR2032 or compatible) battery is installed correctly on the controller card and the protective label removed.

Software Installation

- Promise SuperTrak66 BIOS is loading correctly at system boot
- At least one array has been created with the drive(s) attached to the SuperTrak66 using SuperBuild or SuperCheck
- If the array is to be a system drive, it is partitioned & formatted appropriately
- SuperTrak66 Message Agent is installed (automatically installs I/O driver component)
- If this system is to be a monitoring gateway or server, the SuperTrak66 Message Server is installed
- If this system is to be used for monitoring SuperTrak66 server systems, SuperTrak66 SuperCheck utility is installed.

2.3 Hardware Installation

- 1) Install EDO memory into the SIMM slot at SIMM1 (see Figure 1). The minimum memory requirement is 8MB. For optimal performance, Promise recommends at least 16MB. Some units may ship with 16MB memory.
- 2) Remove protective label for 3V (CR2032 or compatible) battery shipped with the card (see Figure 1) that supports NVRAM.
- 3) Remove protective label for audible alarm buzzer (see Figure 1).

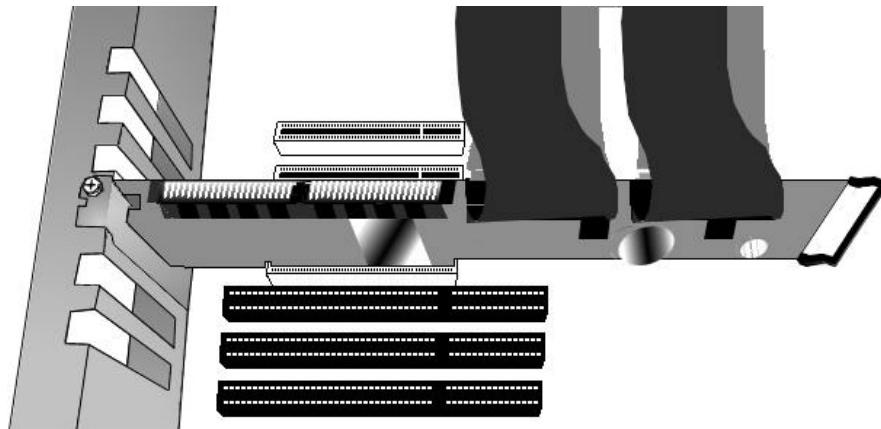


Figure 2: Typical controller card installation

- 4) Install the controller card into an available PCI slot on the system motherboard and secure it in place.
- 5) Set the hard drives to "Master", "Cable-Select", or "Single" (depending on drive specification).
- 6) Next, install hard drives into the PC chassis. Attach black end of Ultra ATA/66 cable and power cable to back of each drive..
- 7) Connect the blue end of the Ultra ATA/66 cable to any IDE channel available on the SuperTrak66.

NOTE: Device support for drives is provided by the SuperTrak66 system. No changes are necessary in the motherboard CMOS/BIOS Setup for resources or drive types regarding hard drive configuration. Changes may or may not be necessary for PCI device setup.

2.4 Software Installation

This section covers various software-related installation issues for SuperTrak66. These include: BIOS Utility, Network Connection, Driver Installation, and Application Installation.

BIOS Utility

The SuperTrak66 BIOS utility ("SuperBuild") must be used to create the first array(s) on the SuperTrak66. Arrays may subsequently be created through the SuperCheck utility. You must have an array on SuperTrak66 in order to use the drives for data storage. See Chapter 3 for more information about using the SuperBuild BIOS utility to create an array.

Network Connections

SuperTrak66 may be monitored over a network. The following topics require attention during the configuration. If you will not be monitoring SuperTrak66 via a network, you may skip this particular section. *Please note that this is NOT intended as a thorough networking/troubleshooting reference for LAN or WANs.*

Firewall

For remote connection to a LAN on which a SuperTrak66 Message Server is installed, it is very possible that you may be required to communicate through a firewall. To do this, you must have a communication port open for the IP address on which the Message Server resides, or the IP address must be fully open for communication.

The network administrator should either: a) open port 711 (decimal) on the IP address; or, b) open all ports for the IP address. This will allow the SuperCheck software to initiate a TCP/IP connection from the remote location through the firewall port 711 to the Message Server located at the destination IP.

TCP/IP

Proper TCP/IP setup is crucial for the SuperTrak66 software package to function correctly. Without TCP/IP installed and configured correctly, the software will not be able to make connections to the SuperTrak66 RAID server(s) for remote access and maintenance.

For your reference convenience, we have included an outline of the Windows NT TCP/IP installation procedure:

- 1) Open the Control Panel under "Settings.." from the Start Menu.
- 2) Double-click the "Network" icon.
- 3) Select Protocol Tab.
- 4) Press "Add" button.

- 5) Highlight “TCP/IP Protocol” from the listbox, then click on the “OK” button.
- 6) Windows NT will proceed to install the TCP/IP protocol stack.
- 7) Click on the “Configure” button.
- 8) Enter the network IP address. If you do not know the IP address, contact your network administrator. An incorrectly configured IP may result in network conflicts.
- 9) Enter the Subnet Mask (generally 255.255.255.0, but contact your network administrator if you are uncertain.)
- 10) Save all the settings, shut down, and restart the server.
- 11) You may test the server response by issuing a “ping” command from the command prompt shell on another workstation on the LAN. (“ping xxx.xxx.xxx.xxx” where “xxx” corresponds to the IP address you configured the server for). A reply will be received if both systems are configured correctly.

NOTE: For further information on IP addressing, and TCP/IP configuration, consult the Windows NT documentation.

Driver Installation for Existing Windows NT 4.0 System

The following procedure details the installation of the SuperTrak66 drivers on a system with Windows NT version 4.0 or later already installed. Under Windows NT 4.0, insert the included "SuperTrak66 Driver" diskette in drive A: when prompted for the drivers.

- 1) From the "Start" menu, choose "Control Panel" in the "Settings."
- 2) In the "Control Panel" window, double-click on the "SCSI Adapters" icon.
- 3) In the "SCSI Adapters" window, choose the "Drivers" tab.
- 4) Press the "Add" button.
- 5) Choose "Have Disk" in the "Install Driver" window (insert Driver diskette)
- 6) Enter the directory location of the driver in the text box (A: or use browse button to locate file).
- 7) Click on the filename "sptrak.inf" and press "Open" button.
- 8) Select "Promise Technology, Inc., SuperTrak IDE Controller" and press "OK."
- 9) "Select SCSI Adapter Option" will be displayed. Press "Install," which initiates the installation operation.
- 10) After a successful installation, the "SCSI Adapter Setup" dialog box will appear which states that the, "Promise Technology, Inc., SuperTrak IDE Controller" has been installed.
- 11) Restart the system.

To check if the driver installation was successful, perform the following steps:

- 1) After restarting the system, return to the Control Panel.
- 2) Click on the "SCSI Adapters" icon.
- 3) Once in the "SCSI Adapters" window, click on the "Devices" tab.
- 4) In the box titled "SCSI Adapters and connected devices are listed below," the Promise Technology, Inc. SuperTrak IDE Controller" should appear.

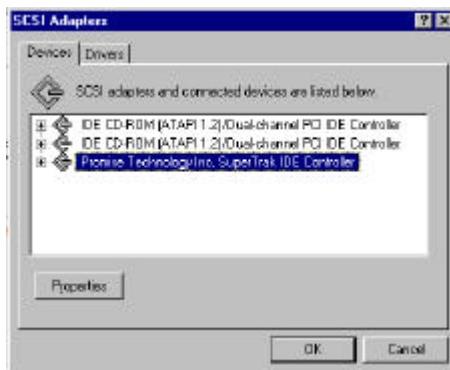


Figure 3: Driver Installed Screen in Existing NT

Driver Installation During New Install of Windows NT 4.0

Under Windows NT 4.0, insert the included "SuperTrak66 Driver" diskette in drive A: when prompted for the drivers.

- 1a) Floppy Install: Boot the computer with the Windows NT installation diskettes.
 - 1b) Floppyless Install: Boot from floppy and type "WINNT /B". After files have been copied, the system will reboot. On the reboot, press <F6> when the message "Setup is inspecting your computer's hardware configuration..." appears.
 - 1c) CD-ROM Install: Boot from the CD-ROM and press <F6> when the message "Setup is inspecting your computer's hardware configuration..." .
 - 2) When the "Welcome to Setup" window is generated, press "Enter."
 - 3) In the "Setup Method" dialog box, press "Enter" to confirm "Express Setup."
 - 4) Press "S" to add the SuperTrak66 controller.
 - 5) Select "Other (requires disk provided by hardware manufacturer)."
 - 6) Specify "a:\", insert the Promise driver diskette into drive A: and press ENTER.
 - 7) Select the "Promise Technology Inc. SuperTrak Controller" and click "OK".
- NOTE:**for CD installations, you must also specify the driver for your CD-ROM adapter (i.e. if using an ATAPI CD-ROM, specify the IDE 2.1 controller)
- 8) Follow the normal setup installation procedure.

To check if the driver installation was successful, perform the following steps:

- 1) After restarting the system, return to the Control Panel.
- 2) Click on the "SCSI Adapters" icon.
- 3) Once in the "SCSI Adapters" window, click on the "Devices" tab.
- 4) In the box titled "SCSI Adapters and connected devices are listed below," the Promise Technology, Inc. SuperTrak IDE Controller" should appear.



Figure 4: Driver Installed Screen in New NT

SuperCheck Installation

There are several layers of the SuperCheck software installation which must be installed on the appropriate network servers/workstations. Make sure the SuperTrak66 I/O device driver is installed as detailed on p. 9 before moving on to install the SuperCheck utility software. Chapter 4 covers the complete SuperCheck installation and reference for network and stand-alone systems.

NOTE: *It is important to read over the Chapter 3 introduction before skipping ahead to Chapter 4. The introduction contains important initialization information.*

NOTES

SuperBuild™

This chapter contains a *Setup Task Quick Reference* which is a checklist of tasks needed to initialize your SuperTrak66 system, followed by a step-by-step instructional breakdown of each task. The second half of this chapter is a visual and technical description reference including software screen shots.

3.1 Setup Task Quick Reference

Use this reference to make sure you have completed all the tasks necessary to get your SuperTrak66 server running. See Chapter 4.0 for visual references and technical descriptions of the SuperBuild™utility. See Chapter 6 for more information on RAID arrays (what types of arrays to configure for what applications, etc.).

NOTE: *The following checklist is primarily intended for those who require SuperTrak to provide support for an array as a system or bootable drive. If the array is only to be used for data storage, then this checklist should be skipped. Use the SuperCheck utility to create and manage your arrays. It is recommended that you still read over the “Common tasks” (below) to familiarize yourself with the functionality of the BIOS SuperBuild utility.*

Setup Checklist

- Power up the system and watch for the Promise SuperTrak66 BIOS messages.
- Enter the SuperBuild utility (press Ctrl-F).
- For easiest configuration, use the “Auto setup” to create a single array using all drives attached to the SuperTrak66; otherwise, use the “View/Define array” menu to create a new, custom array.
- If you are defining more than one array and booting from one of these arrays, make sure to select the appropriate array as the “boot array” from the “View/Define array” menu. Use the up/down arrow keys to highlight the array and press the spacebar.
- Restart the system for the changes to take effect.
- Watch for the Promise SuperTrak66 BIOS messages, ensuring that the messages reflect the array(s) you have created.

Common Tasks

The following portion of the “Quick reference setup” is a listing of steps necessary to complete each type of task required for setting up the SuperTrak66 system through the SuperBuild BIOS utility. The tasks are not intended to be a complete, step-by-step guide to setting up the SuperTrak66 system successfully.

NOTE: *There are some advanced RAID array configuration options, such as cache settings, which require the SuperCheck utility to set up.*

Entering the BIOS SuperBuild utility

When the system boots, you will see the Promise SuperTrak66 BIOS sign-on banner appear on the screen. Once the BIOS identifies arrays attached to the SuperTrak66 controller, it will offer an option to enter the SuperBuild utility. Press **[Ctrl-F]** to initialize SuperBuild and display the Main Menu.

Creating a disk array from the Auto Setup screen

- 1) From the Main Menu press **[1]** to display the Auto Setup screen.
- 2) Choose an option in the “Optimize array for:” field. The choices offered are RAID 0 (striping), RAID 1 (mirroring), RAID3 (striping w/ dedicated parity), and RAID5 (data striping w/ striped parity), and JBOD (spanning array).
- 3) Press the **[Ctrl-Y]** keys. SuperBuild will build an array automatically from the hard drives connected to the SuperTrak66 controller.

Creating an array from the View / Define Array screen

- 1) From the Main Menu press **[3]** to display the View/Define array screen.
- 2) Select the array number you want to define using the arrow keys, and press **[Enter]**. This will generate the “Define array menu/View drive assignments” screen.
- 3) Highlight the “RAID Level” column using the Tab key, and then cycle through the choices using the spacebar to select the RAID level you wish to use for the new array.
- 4) Using the down arrow, select from the available drives (“Drive assignments” section) to assign them as members of the new array. Use the spacebar to toggle between Yes or No or press **[Y]** under the “Assignment” column to assign the drive; **[N]** unassigns the drive.
- 5) When you are finished assigning drives as array members, press the **[Ctrl-Y]** keys to save the array member assignment. The utility will return to the “Define array” menu and the new array will now be visible in the list.

Deleting an array from the Delete Array screen

- 1) From the Main Menu press **[4]** to display the Delete Array screen.
- 2) To delete an array, highlight the array you wish to delete using the up/down arrow keys.
- 3) Press the **[Delete]** key.
- 4) Press the **[Ctrl-Y]** keys to confirm deletion.

Setting the Halt On Error option

- 1) From the Main Menu press **[5]** to display the Controller Configuration screen.
- 2) Use the Space Bar to toggle the Enable/Disable option of the Halt On Error feature.

3.2 SuperTrak BIOS Reference

When the SuperTrak66 BIOS loads during system boot time, it displays pertinent information about the RAID arrays which it finds, then displays “Press <Ctrl-F> to enter Array BIOS Configuration Utility” (see below).

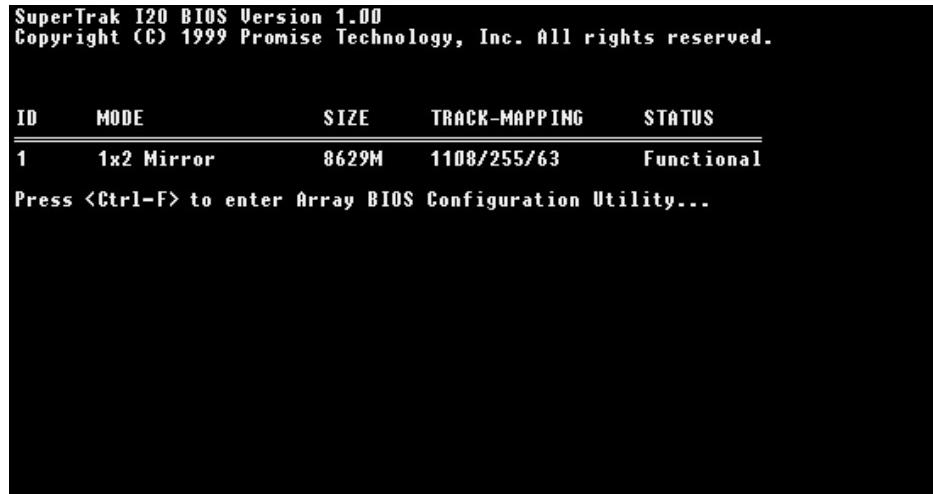


Figure 5: SuperTrak66 BIOS startup

The information displayed in the form of a small table contains the following properties:

ID

These represent a unique ID number assigned to each RAID array identified by the BIOS.

Mode

This identifies the RAID mode configuration that the array is defined as.

Size

This is the MB (Megabyte) data size of the indicated array.

Track-Mapping

This is the CHS (Cylinder/Head/Sector) equivalent of the array geometry as hosted by the SuperTrak66 BIOS int 13h disk services.

Status

Shows one of three possible array status conditions:

Functional

The array is fully operational, and no problems are present.

Critical

The array is operational, but has lost its fault tolerance. For RAID array levels 1, 3 and, 5 the array contains a failed drive. The user should identify and replace the failed drive.

Offline

The array is no longer operational and the SuperBuild utility cannot rebuild the array. The array must be rebuilt from the last tape backup or other device. The user should identify and replace the failed drives. For RAID levels 1, 3, and 5, at least two or more drives in the array have failed. For a RAID 0 array, at least one drive has failed.

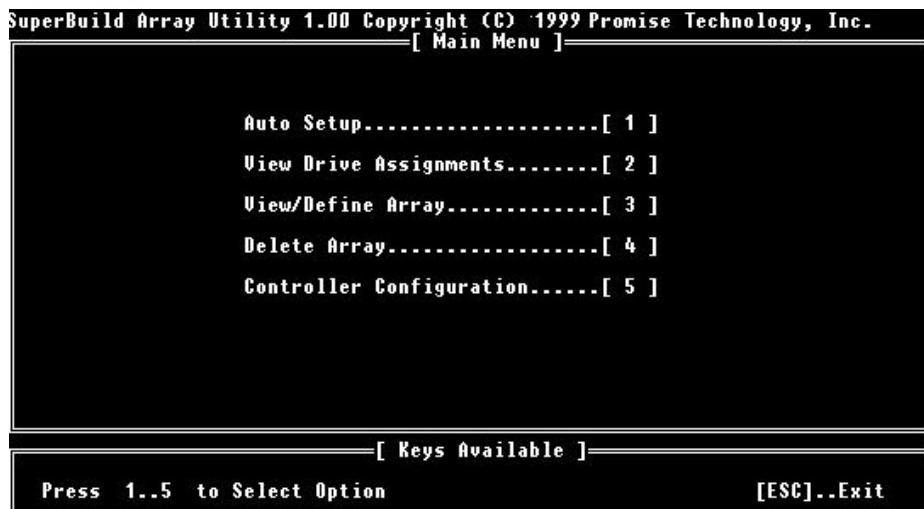
Main Menu

Figure 6: SuperBuild Array Utility Main Menu

Note: *This menu uses number keys 1 through 5 to select options. Many of the subsequent menu windows use: arrowpad keys to highlight options. Use the spacebar to cycle through choices for the selected option. Enter (return) to select a highlighted option. Escape to exit the menu.*

3.3 Auto Setup

By pressing the [1] key while in the Main Menu screen, the “Auto Setup” screen is displayed (see below). It is divided into sections: “Auto setup options menu”, “Auto setup configuration,” and “Keys Available”.

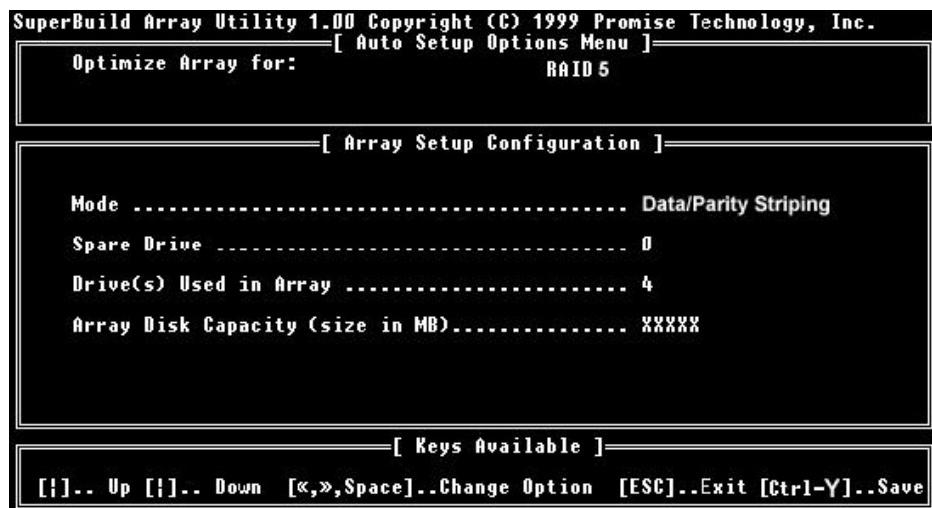


Figure 7: Auto setup screen

Auto Setup Options Menu

This section of the screen is the only selectable portion. Depending on the number of drives recognized by SuperTrak66, the Auto Setup Options menu may limit the number of array types available. You choose a RAID level in the “Optimize array for:” field. The choices that may be offered are: RAID 0 (striping), RAID 1 (mirroring), RAID3 (striped data + parity drive), RAID5 (data and parity striping), and JBOD (spanning).

Optimize Array for:

Selections in this field correspond to the items found in the Mode field in the Array Setup Configuration section (see table below).

RAID Level	Mode
RAID 0	Stripe
RAID 1	Mirror (2 drives)
RAID 0+1	Stripe + Mirror (4 drives)
JBOD	Spanning
RAID 3 (needs 3 drives min.)	Stripe Data + Parity

RAID 5 (needs 3 drives min.)	Data/Parity Striping
------------------------------	----------------------

Auto Setup Configuration

This section of the Auto Setup Menu contains four fields:

Mode

Items in this field correspond to the selections found in the “Optimize Array For” field in the Auto Setup Menu section (see table on page 19).

Spare Drive

Auto setup does not allow configuring a hot spare drive. If you wish to have a spare drive, you must use the View/Define Array Menu (see page 22) and manually assign the exact number of drives you want included. Any attached drive not configured in a fault tolerant array will be recognized as a spare drive and will be automatically added to the array when a failed drive is detected. An automatic data rebuild will be performed to restore fault tolerance as quickly as possible. At a later time, the failed drive can be physically removed and an extra drive added in its place to function as the new spare drive.

Drive(s) Used In Array

This field displays the number of drives used in the array. This does not include the spare drive(s).

Array Disk Capacity

This field displays the capacity available in the entire array in MB (megabytes), not counting spare drives.

When you have completed your selections, press the **[Ctrl-Y]** keys. SuperBuild will create an array automatically from the hard drives connected to the SuperTrak66 controller. SuperBuild will then ask to reboot the system.

3.4 View Drive Assignments

By pressing [2] on the "Main Menu" screen, the "View Drive Assignments" screen is displayed (below). This screen does not allow modifications to any of the four fields.



Figure 8: View Drive Assignments screen

Channel: ID

This field shows the SuperTrak66 controller channel ID (1-4) to which a particular drive is attached.

Drive Model

This field identifies the manufacturer, model, and model number (if applicable) of each drive installed on the SuperTrak66 card.

Capacity (MB)

This field reflects the capacity in MB (megabytes) of each drive.

Assignment

This field identifies the array to which each drive belongs. Assigned drives are labeled with their disk array (1,2,3, etc...). Unassigned drives are labeled "free," and are considered "hot spares" to provide fault tolerance. Unassigned drives may also be used to create a new array at any time.

3.5 View / Define Array

By pressing [3] from the “Main Menu” screen, the “Define Array Menu” is displayed (see Figure 9 below). There may already be existing arrays configured on this screen or it may appear as indicated in the figure below. The Define Array screen allows users to manually begin the process of defining both drive elements and RAID levels for each disk array. Users will commonly create one or two drive arrays with SuperTrak66.

To manually create arrays from the Define Array Menu below, use the arrow keys to highlight an array number you wish to define, and press [Enter] to select. The Define Array Definition Menu will next appear that allows drive assignments to the disk array (see p.23).

NOTE: Promise recommends that if you are creating a single array using SuperTrak without a "hot" spare drive, do so through the “Auto setup” screen under selection [1] of the Main Menu.



Figure 9: Define Array Menu screen

Array No.

The Array Number column displays all arrays (defined or undefined) addressable by the SuperTrak66.

RAID Level

This column indicates the RAID type applied to a particular array.

Total Drv

This column displays the number of drives configured for array.

Capacity (MB)

The values in this column reflect the total capacity in MB (megabytes) for that array.

Status

This column displays the state information for each array. Status definitions can be found at the beginning of Section 3.2.

Define Array Definition Menu

Selecting an Array # from the Define Array Menu brings up the Define Array Definition Menu screen (see Figure 10 below).

You must first assign the type of RAID Level you want. In the Definition Menu section, use the [Space] key to cycle through array types: RAID 0 (Striping), RAID 1 (Mirroring), RAID 3 (Stripe Data + Parity), RAID 5 (Data/Parity Striping) or JBOD (Spanning). See Chapter 5 about RAID levels.

NOTE: Promise recommends the use of identical drives when creating a new array for best results.

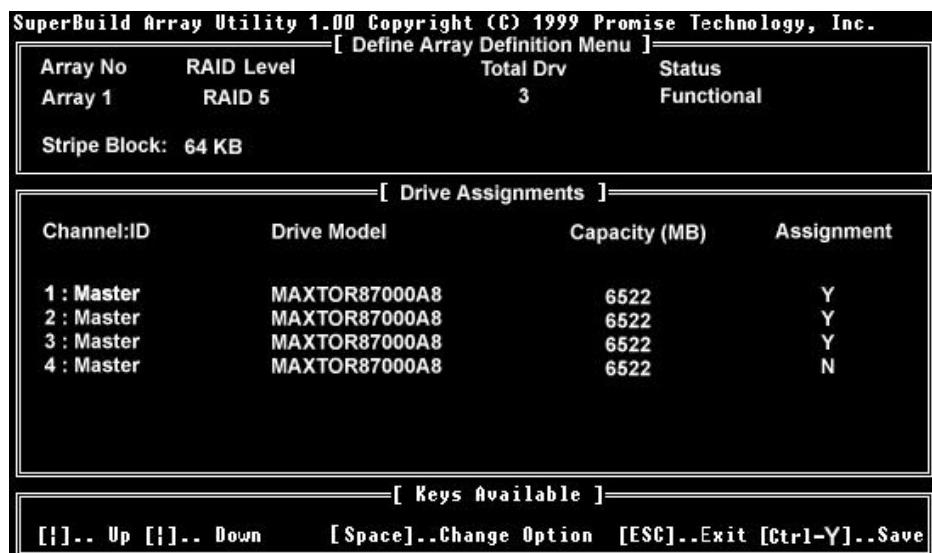


Figure 10: Define Array Definition Menu screen

Creating a RAID 0 Array

In the Definition Menu section (see Figure 10), use the [Space] key to cycle through array types and select "RAID 0" for RAID Level. You will be joining the drives you assign to this array together and splitting (or striping) data writes among the members of the array. Each member of such an array will no longer be readable as an independent drive if removed from the array. The capacity of the array will equal the capacity of the smallest drive multiplied by the number of drives within the array.

Changing Stripe Block Size

For RAID 0 arrays, SuperTrak66 will allow you to change the stripe block size. The default is 64K. However, certain application environments may perform optimally with different values (see drive manufacturer for additional information on block sizes). To change the block size, use the arrow keys to highlight the field and the spacebar to cycle through the available size selections.

Assigning Drives to a RAID 0 Array

Under the [Drive Assignments] section, highlight a drive and, with the [Space] bar, change the Assignable option to "Y" to add the drive to the disk array. Press <Ctrl-Y> to save the disk array information. The initial Define Array Menu screen will appear after you have pressed [Ctrl-Y]. You will see the array defined. You may select another Array number to define OR you may ESC to exit and return to the Main Menu of SuperBuild.

Creating a RAID 1 Array

In the Definition Menu section, use the [Space] key to cycle through array types and select "RAID 1" for RAID Level. You will be creating two drives or two pairs of striped drives (4 drives total) with identical data on them. The capacity of the array will be half that of the total number of mirrored drives within the array. No stripe block size option is available when mirroring 2 drives. In a 4drive RAID 0+1 array, you can set the stripe size for the two pairs of striped drives.

If you are creating a Mirroring array containing two drives only, you will use either two brand new drives, or one drive that contains existing data and a second drive that you wish to mirror data. There is an additional window that will appear in order to create such an array.

Creating a Mirroring Array with Two New Drives

As described above, if you select a mirroring array and wish to use two new assigned drives, follow the directions here. Under the [Drive Assignments] section, assign both new drives to the array and save the information with <Ctrl-Y>. The following window will appear in order to create the array. Confirm "N" or "Create Only" to the message shown in order to proceed.

Do you want the disk image to be duplicated to another? (Yes/No)
Y - Create and Duplicate
N - Create Only

Figure 11: Two-Drive Mirroring Dialogue Window

Once the process is complete, you will be returned to the [Define Array] menu where the array will be shown as defined.

NOTE: you may choose to attach an unassigned third drive to SuperTrak66 to act as a "hot spare" replacement (see p. 28).

Creating a Mirroring Array with an Existing Data Drive

SuperBuild will create a mirrored array using an existing drive with data. In the [Drive Assignments] section, you must assign the existing data drive and another drive of the same or larger capacity to the Mirroring array. Press [Ctrl-Y]. Create the array by confirming "Y" or "Create and Duplicate" to the message window that appears (same message as in Figure 10 above):

Do you want the disk image to be duplicated to another? (Yes/No)
Y - Create and Duplicate
N - Create Only

You will be prompted to select the Source drive from the two available drives. Arrow down to the drive that has existing data to be mirrored and press [Enter] to select at the message window below:

[Please Select A Source Disk]		
Channel:ID	Drive Model	Capacity (MB)
1 : Master	MAXTOR87000A8	6522
2 : Master	MAXTOR87000A8	6522

Figure 12: Source Disk Window for Mirroring Existing Data

When prompted, press Y to confirm Yes to start duplicating the existing source data to the target drive. **NOTE:** all target drive data will be erased. Make sure you choose the correct drive.

After the process is complete, SuperBuild will ask you to reboot the system. You may then return to the [Define Array] Menu to create an additional array.

NOTE: you may choose to attach an unassigned third drive to SuperTrak to act as a "hot spare" replacement (see p. 28).

Assigning Four Drives to a Mirroring Array

Under the [Drive Assignments] section, assign all drives to the array and save the information with <Ctrl-Y>. SuperBuild will automatically create two striped pairs of drives (RAID 0+1). You will be returned to the [Define Array] Menu where the array will be shown.

Creating a RAID 3 Array

Use the information in Chapter 6 to help determine the selection of a RAID 3 array. In the [Define Array Definition Menu], use the [Space] key to cycle through array types and select "RAID 3" for RAID Level. You must configure such an array using a minimum of three drives since additional parity information must be calculated and stored on a dedicated "parity" drive.

You will be joining either two (if 3 drives are assigned) or three (if 4 drives are assigned) of the drives together and splitting (or striping) write data among them. Each member of such an array will no longer be readable as an independent drive if removed from the array. The capacity of the array will equal the capacity of the smallest drive multiplied by the total number of drives assigned to the array minus one.

NOTE: If you choose to use just three drives, a fourth unassigned drive can be your "hot" spare drive in case of drive failure (see p. 28).

Changing Stripe Block Size for RAID 3 Array

For RAID 3 arrays, SuperTrak66 will allow you to change the stripe block size. The default is 64K. However, certain application environments may perform optimally with different values (see drive manufacturer for additional information on block sizes). To change the block size, use the arrow keys to highlight the field and the spacebar to cycle through the available size selections.

NOTE: you can not change the block size AFTER the array has been created.

Assigning Drives to a RAID 3 Array

Under the [Drive Assignments] section, highlight a drive and with the [Space] bar change the Assignable option to "Y" to add the drive to the disk array. Press <Ctrl-Y> to save the disk array information. The initial Define Array Menu screen will appear after you have pressed [Ctrl-Y]. You will see the array defined. You may ESC to exit and return to the Main Menu of SuperBuild.

Creating a RAID 5 Array

Use the information in Chapter 6 to help determine the selection of a RAID 5 array. In the Definition Menu section, use the [Space] key to cycle through array types and select "RAID 5" for RAID Mode. You must configure such an array using a minimum of three drives because of parity data generation.

You will be joining either two (if 3 drives are assigned) or three (if 4 drives are assigned) of the drives together and splitting (or striping) data written among them. Each member of such an array will no longer be readable as an independent drive if removed from the array. The capacity of the array will equal the capacity of the smallest drive multiplied by one less than the total number of drives assigned to the array.

NOTE: If you choose to use just three drives, an unassigned fourth drive can be your "hot spare" drive in case of drive failure (see p. 28).

Changing Stripe Block Size for RAID 5 Array

For RAID 5 arrays, SuperTrak66 will allow you to change the stripe block size. The default is 64K. However, certain application environments may perform optimally with different values (see drive manufacturer for additional information on block sizes). To change the block size, use the arrow keys to highlight the field and the spacebar to cycle through the available size selections.

NOTE: you can not change the block size AFTER the array has been created.

Assigning Drives to a RAID 5 Array

Under the [Drive Assignments] section, highlight a drive, and, with the [Space] bar, change the Assignable option to "Y" to add the drive to the disk array. Press <Ctrl-Y> to save the disk array information. The initial Define Array Menu screen will appear after you have pressed [Ctrl-Y]. You will see the array defined. You may ESC to exit and return to the Main Menu of SuperBuild.

Creating a Spanning Array

In the [Define Array] section, use the [Space] key to cycle through array types and select "JBOD" for RAID Level. You will be joining the drives you assign to this array together as a single "virtual" drive grouping. Data will be stored to the first drive until it is physically full. SuperTrak66 then will save data on to the next drive in sequence. Each member of an array will no longer be readable as an independent drive if removed from the array. The capacity of the Spanning array will equal the total capacity of all drives assigned.

NOTE: Since spanning uses the full capacity of each assigned drive, different capacity drives may be used without adversely affecting the performance of the array. No other RAID benefits are provided (see Chapter 6).

Assigning Drives to a Spanning Array

Under the [Drive Assignments] section, highlight a drive, and, with the [Space] bar, change the Assignable option to "Y" to add the drive to the Spanning array. Press <Ctrl-Y> to save the disk array information. The initial Define Array Menu screen will appear after you have pressed [Ctrl-Y]. You will see the array defined. You may press the [ESC] key to exit and return to the Main Menu of SuperBuild.

"Hot" Spare Drives

Any extra drive that is not assigned to a fault tolerant disk array (RAID 1, 3, or 5) will be recognized as a Spare Drive. In Figure 10 on page 23, the fourth drive shown is unassigned and would be used as a hot spare for this RAID 3 array. Such a drive is automatically added to an array once a disk member of the array has been detected as "failed." To restore fault tolerance as quickly as possible, SuperTrak66 begins to perform an automatic data rebuild on the "spare" drive in the background without the need to restart the system. At a later time, the failed drive can be physically removed and an extra drive added in its place to function as the "spare" drive.

Setting a Disk Array As Bootable

You may select which defined array will be used as the bootable array from the [Define Array] menu (Item [3] in the Main Menu). Highlight the desired array and press the [Space] bar key. The system will now recognize this as the first array and an * asterisk will appear next to the array number indicating it as bootable (see Figure 13 below).

NOTE: Your PC or server must be configured to use SuperTrak66 as the bootable controller. The system will then use the bootable array as the (fixed) boot C: drive.

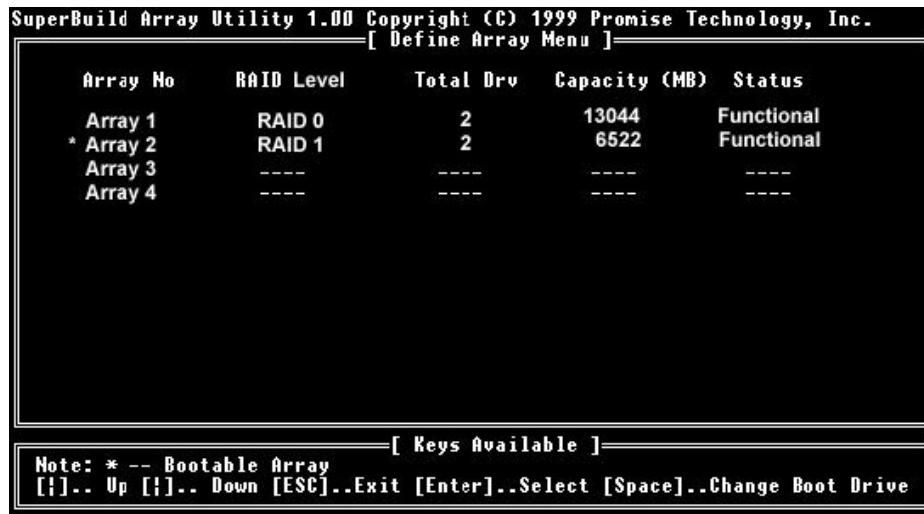


Figure 13: Assigning Bootable Array

Disk Array Recognition Order

During startup, the disk arrays on the SuperTrak66 are recognized in this order:

- 1) The array set to bootable in the SuperBuild™ setup
- 2) By the Array number (i.e. Array 1, Array 2.). This would be involved in determining which drive letters will be assigned to each disk array.

Saved Disk Array Information

NOTE: Promise suggests that users record their disk array information for future reference.

All disk array data is saved within the reserved sector of each arrayed drive.

Another feature of the SuperTrak66 disk array system is to recognize drive members even if drives are moved to different IDE channels on SuperTrak66. Since each drive's array data identifies itself to the array, it is possible to reorder the drives without modifying the array setup. This is valuable when adding drives, or during a rebuild.

3.6 Delete Array

By pressing **[4]** from the "Main Menu" screen, the "Delete Array Menu" screen is displayed (see below).



Figure 14: Delete Array Menu

NOTE: Deletion of an existing disk array could result in the loss of all data from the hard drives of the affected array. Record all array information, including array type, disk members, and stripe block size, in case you wish to undo a deletion. Arrays may possibly be recovered after deletion by immediately re-defining the array with information identical to the original configuration.

To delete an array, highlight the Array you wish to delete and press the [Del] key.

Confirm yes to the following warning (see Figure 15 message using the <Ctrl-Y> keys to continue array delete. Press any other key to Abort the deletion process.

Are you sure you want to delete this array?
Press Ctrl-Y to Delete, others to Abort

Figure 15: Delete Array Dialogue Box

3.7 Controller Configuration

By pressing [5] on the “Main Menu” screen, the “Controller Configuration” screen will appear (see below).

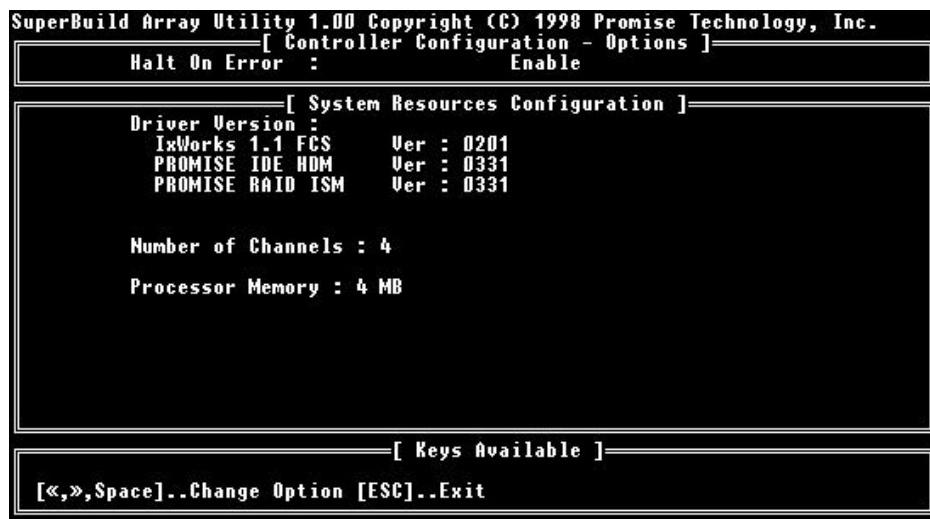


Figure 16: Controller configuration screen

This screen allows you to modify the "Halt On Error" option (enable/disable) by pressing the spacebar. "Halt On Error" is enabled if you want the system to halt processing during bootup if the SuperTrak66 BIOS determines that there is an array error. All other information in the System Resources Configuration section is for display only. The information presented here is as follows:

Driver Version

Shows the driver software version of the IxWorks operating system, Promise Hardware Driver Module (HDM), and Promise Intermediate Software Module (ISM).

Number of Channels

Shows the number of channels attached to the SuperTrak66 controller which may be used to connect one drive each for use in RAID arrays.

Processor Memory

Shows the size of the SIMM plugged into the SuperTrak66 controller card.

SuperCheck™Setup

SuperCheck™is Promise Technology' s array and system monitoring utility. With SuperCheck, you can monitor you're arrayed hard drives, rebuild arrays, and check the operation of server-level components (fans, etc.). There are four basic software components to SuperCheck: the Message Server, the Message Agent, the I₂O driver, and the SuperCheck Monitoring Utility.

Before moving on to the actual installation process, we will describe where SuperCheck components should be installed in a typical network, what hardware they should be associated with, and what function each component performs.

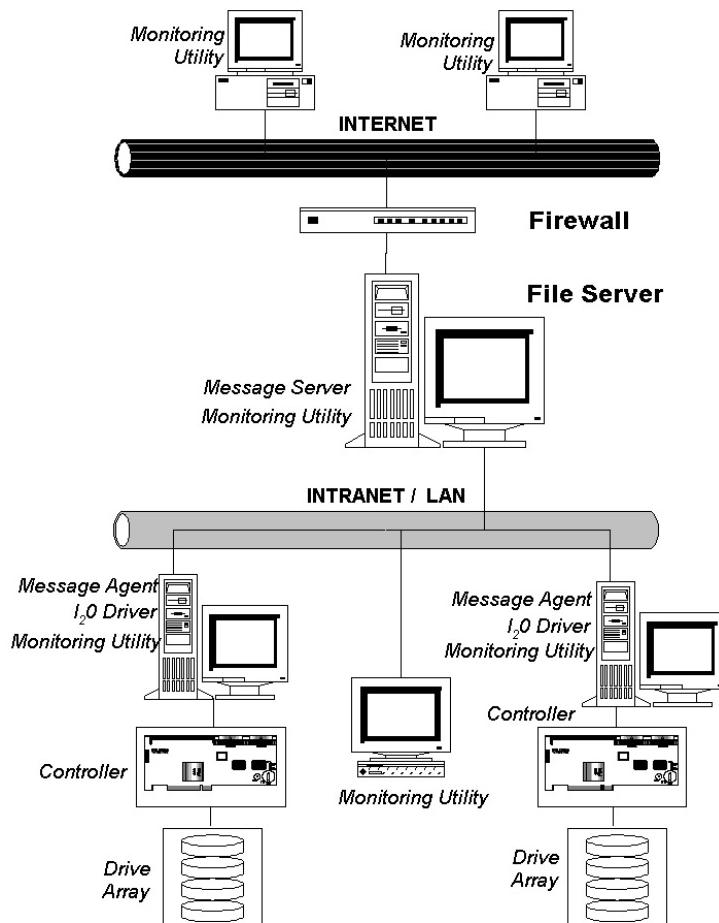


Figure 17: Typical SuperCheck Component Installation

A typical installation of the SuperCheck software components on a network might look like Figure 17 on page 32. The major elements of this network are (from top to bottom):

- Remote Monitoring Stations (via Internet)
- Main Network File Server (with direct Internet access via a Firewall)
- Internal SuperTrak stations (connected via Intranet/LAN)
- Internal Monitoring Station

As shown, different components of SuperCheck are installed on each of the network elements.

- (1) **SuperCheck Utility** - The SuperCheck Monitoring Utility (the GUI that you will be using) may be installed on any system element with a TCP/IP connection where you want to perform monitoring. As shown in Figure 17, the SuperCheck utility is installed on the main File Server, on both SuperTrak66 stations, on two remote systems over the Internet, and on one internal system. *NOTE: For security reasons, array maintenance features can only be performed at the actual server location of a SuperTrak66 array and are unavailable at all other consoles.*
- (2) **Message Server** - The Message Server software component serves as the main link between all SuperTrak arrays residing on the network, the internal Intranet/LAN (via TCP/IP protocols), and remote monitoring stations communicating via the Internet. For remote Internet monitoring, the Message Server component must be installed on a machine with an unique IP address that is visible to the Internet (as in our example). From its central communications position, the Message Server software interfaces with all SuperTrak arrays located on the LAN and allows all system elements (with the GUI installed) to see those arrays. The File Server where the Message Server software is installed may or may not house a SuperTrak array. In our example on Figure 17, it does not. One reason may be that the File Server is not running Windows NT 4.0.
- (3) **Message Agent** - The Message Agent software component must be installed on each file server that will house a SuperTrak array. The Message Agent communicates with both the Message Server and the SuperTrak66 controller. Installing the Message Agent automatically installs the I₂O device driver. *NOTE: All stations containing SuperTrak66 must be running Windows NT 4.0. Maintenance on SuperTrak arrays can be performed only using these stations .*
- (4) **I₂O Device Driver** – Installed automatically with the Message Agent, the driver allows the station to communicate directly between the Promise SuperTrak controller card and Windows NT. This permits Windows NT to recognize RAID arrays attached to it.

4.1 SuperCheck Installation

SuperCheck has many component modules, which, depending on your scenario, may or may not need to be installed on your system. The following “*Installation Scenarios*” describe what components are to be installed, and where.

Installation Scenarios

The following installation scenarios each detail a separate software installation model. Some scenarios are based on a single system, while others involve multiple systems on a LAN or WAN. Select the scenario(s) which best match your needs, then follow the installation steps listed for that scenario. Each of these components may be installed by following the “Common Components” installation procedure outlined on p.36. Proceed to section 4.2 once installation is complete to begin “*SuperCheck Administration*.”

NOTE: All systems with the SuperTrak66 controller will require the Message Agent software (includes automatic installation of the I₂O device driver). Additionally, the component installations that follow require TCP/IP to be already properly installed on the server system. Chapter 2 contains some information about installing TCP/IP on Windows NT.

NOTE: While you may install multiple systems with the Message Server software, Promise recommends against it because the same SuperTrak arrays will be visible from every instance of the Promise Message Server, causing confusion.

All Scenarios:

- 1) If the system is on a LAN or WAN, use the IP address for which your TCP/IP is already configured to establish connections. Otherwise, use IP address 127.0.0.1 for all your TCP/IP configurations; This address will allow all TCP/IP requests to remain within the system.
- 2) Install the Message Agent on the station with SuperTrak66 (automatically installs I₂O driver component).
- 3) Reboot the system AFTER all software components have been installed in order to activate any installed NT services.

To access a SuperTrak system locally

- 1) Install the Message Server on to this system.
- 2) Install the Monitoring Utility onto this system.

To access a SuperTrak system remotely across a LAN

- 1) Install the Message Server software on to one (or more) systems on the LAN. The system(s) chosen for the Message Server may also include a SuperTrak66 station, be an independent station on the LAN, or act as a network server for the LAN. All Message Agents connected on this LAN will be visible from each instance of the Message Server.
- 2) Install the SuperCheck Monitoring Utility software on to each LAN station from which you wish to monitor the SuperTrak66 array(s). When you begin SuperCheck, you will need to perform "Adding a New Message Server" for each Message Server installed in step 1 that you want to monitor.

To access a SuperTrak system remotely on a WAN

- 1) If you have not done so already, install a Promise Message Server on a system on the LAN. The system chosen for the Promise Message Server installation may house a SuperTrak66, be an independent station on the LAN, or house a network server for the LAN. For remote Internet monitoring, the Message Server component MUST be installed on a machine with an unique IP address that is visible to the Internet. All Message Agents connected on this LAN will be visible for each instance of the Promise Message Server.
- 2) Install the SuperCheck Monitoring utility on to each remote station where the SuperTrak system will be monitored. When you begin SuperCheck, you will need to "Add a New Message Server" for the Promise server installed in Step 1 that you wish to access.
- 3) If the WAN to which the Monitoring utility will be connected has any firewall protection measures, you must open the IP address for the system (or systems) containing the Promise Message Server software in order to connect to them over Internet/WAN. TCP/IP port number 710 and 711 are used by the Message Agent to accept in-bound connections.

Common Component Installation

This section outlines the installation procedure for a component of the SuperCheck software package. See the installation scenarios outlined in section 4.1.1 to determine which components to install.

NOTE: If you are re-installing SuperCheck, you must first stop the services for the Message Server and Message Agent for installation to work.

For each component to be installed, do the following:

- 1) Insert the SuperCheck Install Diskette Number 1 into the floppy drive.
- 2) From Windows NT, go to the Start button, and click on “Run...” from the Start Menu.
- 3) Type in:

A:\setup.exe

- 4) The main SuperCheck component installation menu will come up as seen below.



Figure 18: SuperCheck component installation menu

- 5) Select a software component from the menu by clicking it.

- 6) When the “Choose Destination Location” window (see Figure 19 below) comes up, choose a directory to install this component.
- 7) Click “Browse” to select a Destination folder on your system other than the folder suggested by default. Click the “Next” button to accept the location selected.

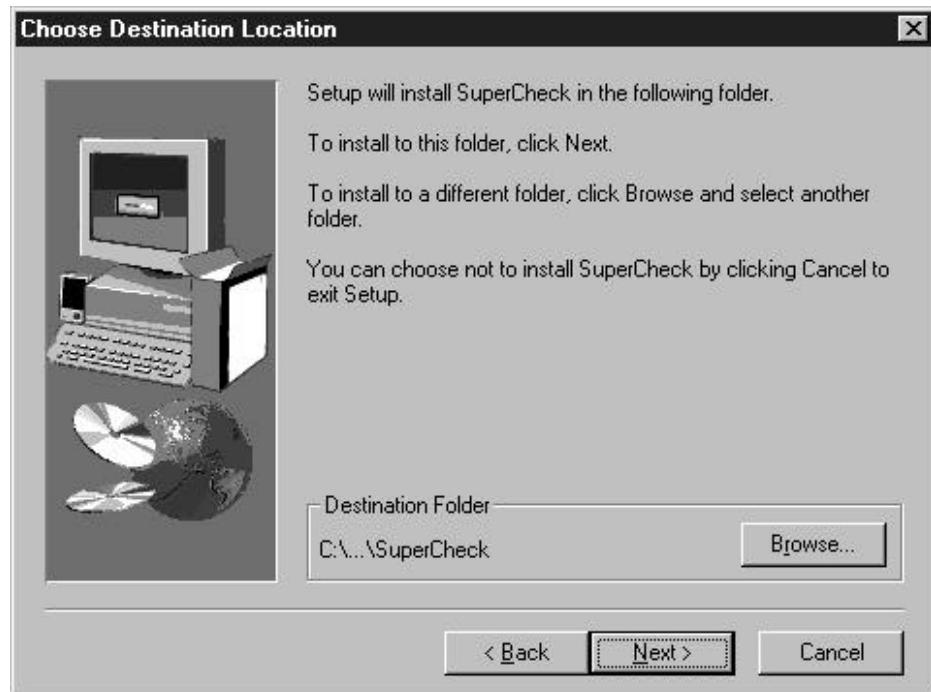


Figure 19: “Choose Destination Location” window

- 8) The “Select Program Folder” window comes up next, shown in Figure 20 below. Choose a Start menu folder to list this item under by entering a folder name, and selecting a “parent” folder from the list under which this new folder’s contents will reside. Click the “Next” button to proceed.

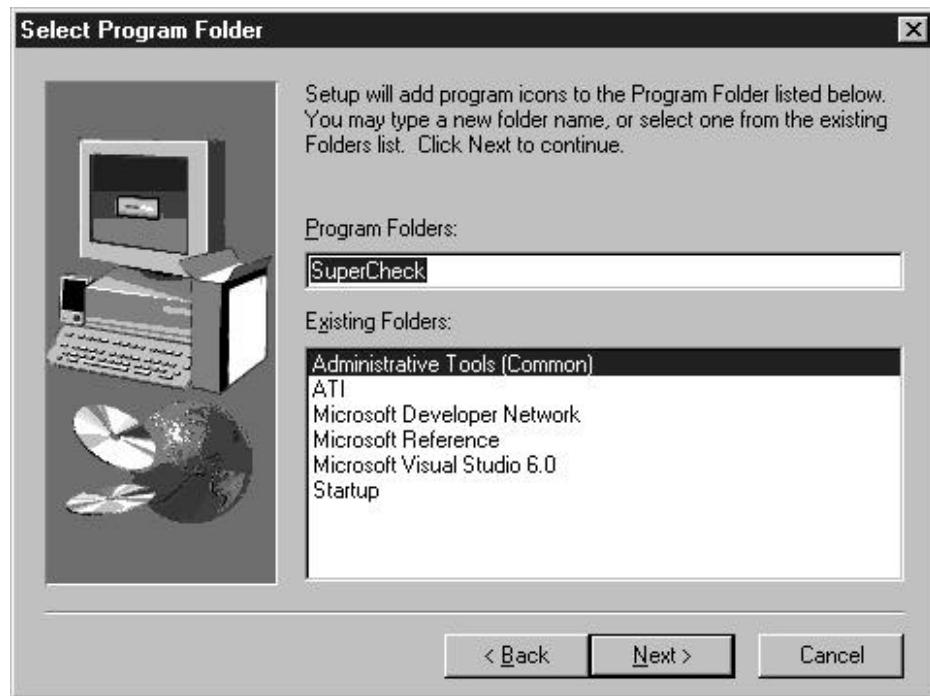


Figure 20: “Select Program Folder” window

- 9) The “Start Copying Files” window will then be displayed (see Figure 21 below). Verify that the proper component(s), destination folder, and program folder you selected for installation are correct, then click on “Next” to proceed. If otherwise, click the “Back” button to go back and make changes.

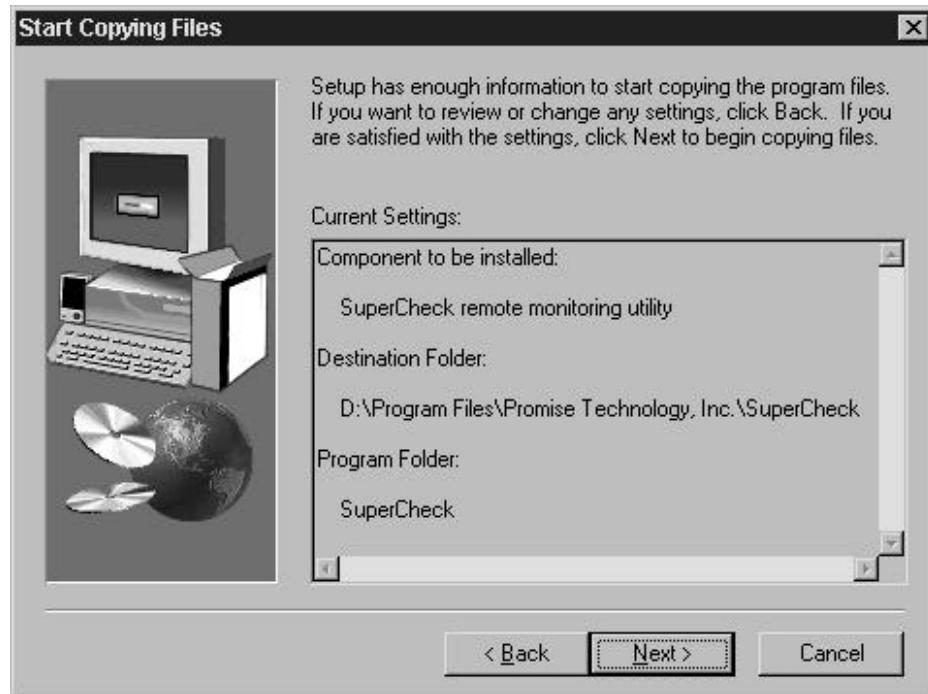


Figure 21: “Start Copying Files” window

- 10) When you see the “Setup Needs the Next Disk” pop-up window (see Figure 22 below), insert the SuperCheck install diskette number 2 in the floppy drive. Click on the “OK” button when it is ready. [You may use the “Browse” button if your installation software resides anywhere other than the floppy diskette to choose the correct source location.]



Figure 22: "Setup Needs the Next Disk" pop-up window

- 11) When the installation is complete, you will see the "Setup Complete" window. Click on the "Finish" button to complete the install process for this component and restart your system.

4.2 SuperCheck Administration

Once installation is complete, you may begin using the SuperCheck utility. There are still a few things that need to be configured before your SuperCheck install and administration are complete. This section will outline tasks necessary to complete SuperCheck Administration. See section 4.3, “*SuperCheck Task Quick Reference*” for details on executing each task. Please refer to Chapter 5 for a complete visual technical reference of the SuperCheck utility.

After you have reviewed and implemented each task below, your SuperCheck installation will be complete. At this point, you will be ready to use your SuperTrak66 RAID server for workload I/Os. Configure your array(s) for operation by partitioning and formatting them under the standard Windows NT “*Disk Administrator*.”

Creating Message Servers & Agents

The first time you run SuperCheck, there will be no Message Server, Message Agent, or array attachments. To establish the server connections, you must determine the TCP/IP address of the system that is running the SuperTrak66 Message Server. You may need your LAN administrator to provide this information. See Section 4.1 for installation scenarios to determine where the Message Server should be installed. Once you have identified the Message Server, perform the “*Create a new message server*” task as detailed in the SuperCheck Task Quick Reference, based on the Message Server system’s IP information.

Once the Message Server has been connected successfully, double-click on the Message Server icon to expand the hierarchical tree to show all attached Message Agents. Log into an attached Message Agent for access to the particular SuperTrak66 system under the Agent’s control.

Adding/Deleting New Users

If additional users require access to a SuperTrak66 Message Agent, the system administrator must create a “new user” under the Message Agent. Each user *may* be given access to various levels of control by an admin. A “*User / Administrator*” may create/delete other user accounts. “*Array Administration*” and “*Adapter Maintenance Rights*” boxes allow access to specific functions of the system.

NOTE: Don’t forget to set a password for any user who has anything more than viewing privileges on a secure network.

Setting Up A New Array

Setting up an array is considered a simple task. Consideration should be given as to the type and configuration of the array. Please see Chapter 6 for RAID configuration ideas and then perform the task of creating the new array (see p. 16)

Array Administration

Array administration is the ability to create, view, modify and delete arrays. The ability to view the status of an array through the utility allows identification of problems which may require user intervention to prevent the array from going down or to restore fault tolerance.

NOTE: *Actual maintenance of the array (deleting, rebuilding, creating, etc.) may be performed only from the system housing the SuperTrak array for security reasons. SuperTrak greys out the maintenance functions for monitoring stations.*

The I/O information (IOP or Controller) subpanes (see page 84) of SuperCheck allow you to see that all drives are functioning correctly. If any drive has failed, SuperCheck will indicate which drive(s) has failed. You also can see the array status to determine if it is functioning, or if it requires critical attention, or if it has begun an automatic rebuild, etc. The array page allows modification of the cache settings for fine-tuning your server's performance. The enclosure pane allows you to verify that the drives are operating within temperature specification, and that the enclosure cooling fans are operational.

NOTE: *Never delete an array unless you have backed up all data which is needed on the array. It may be possible to restore access to an array after the deletion by re-creating the array exactly as it was defined previously. Though the product allows for this procedure, Promise does not recommend it.*

Changing the cache settings will directly affect performance of an array. One way of determining the cache settings which are appropriate for your SuperTrak66 is to use a benchmarking utility. Such a utility stresses the array in a manner consistent with expected typical workload for the array. Choose a configuration that yields the best performance results. It is difficult to make any solid recommendations for cache settings since the array workload will differ from system to system depending on its application.

Setting Up Email Notification

SuperCheck allows the administrator or user to check on array status at anytime. You can also set up email notification so that SuperCheck will notify you of any problems via email messaging services. The notification feature is recommended for any type of array which is to be accessible over a network.

NOTE: *E-mail alert notification is Message Agent-based, not array-based. You will receive email notification for problems related to all arrays connected under the Message Agent.*

The e-mail alert notification may be adjusted for each Message Agent to only send messages for a particular list of events. The administrator can add or remove events from the e-mail alert notification as desired.

Identifying Problems

There are a variety of problems which can be detected by SuperCheck, allowing you to possibly prevent a system crash or data loss due to a hardware problem. The software can determine if an array member has failed, show whether the housing temperature is within operational parameters, or identify a problem with the housing cooling fans. If you experience any failures, please see the Troubleshooting Guide in Chapter 7 for a recommended course of action.

4.3 SuperCheck Quick Reference

The following section is a listing of initial steps necessary to complete each task during administration of a SuperTrak66 server system using the SuperCheck utility. These tasks are not listed in any particular order, nor is this intended to be a complete, step-by-step guide to setting up the software successfully. Please refer to the previous section 4.2, "SuperCheck Administration," for details on what tasks to execute for basic operations."

To Create A New Message Server

- 1) Right-click on the "MyConsole" icon  located in the Tree View area of the SuperCheck main menu (or left-click "MyConsole" from the pull-down selections at the top of the main menu).
- 2) Click on "New", then "Server" in the expanded pop-up menu
- 3) In the message server Information View, enter "Label," and "IP address" information. If you are not sure of the server name, try just entering the IP address.
- 4) Click the "Confirm" button. Once a new Message Server is created, it will create a Message Agent icon  automatically.

Creating A New User

- 1) Log into the Message Agent for which you would like to add a New User.
- 2) Right-click the "Users" icon  in the Tree View area of SuperCheck Main Menu. Then left-click on "New" and select "User" from the pop-up selection.
- 3) A "New User" icon  will appear. Double-click on it to view the user Information View.
- 4) Modify all the user information according to the desired access level for the new user.
- 5) If the user has any level of administrative control, and the server is connected to a WAN, make sure to set a password.
- 6) Click the "Commit" button to update the user information.

Deleting an Existing User



- 1) Select the icon representing the user you wish to delete.
- 2) Right-click the icon and select “Delete” from the pop-up menu.
- 3) Choose “Yes” if this is a user you wish to delete.

NOTE: SuperCheck will not allow deletion of the last account with admin privileges. This protects the admin from being locked out of the system and having to re-install SuperCheck.

Accessing Message Agent



- 1) When clicking on any Message Agent icon (if you are not already logged into the Agent), you will be prompted for password entry.
- 2) Enter the account/password for this system to access the Agent and to gain access to the SuperTrak system.
- 3) Checking the “Save password” box will remember the password for this Message Agent. Future accesses to this Message Agent from this terminal will not require password input.
- 4) Click “OK” to submit the password and access the Message Agent.

Creating a New Array



- 1) Select the IOPx (or Controller) icon on which you want to create an array.
- 2) Right-click the icon. Click “New”, then select “Array” from the pop-up menu.
- 3) Enter the pertinent information in the “Name,” RAID Level” and “Block Size” fields. (Available RAID level selections are based on number of drives to be used; Block Size is selectable from a pull-down list or use the default 64K)
- 4) A new array icon will be created in the Tree View, titled with the text entered into the “Name” field.
- 5) Drag and drop any unassigned drive icon on top of the array icon which you have just created in order to add the drive to the array (unassigned drive icons do not have a red arrow in the upper left corner). Hard drive(s) will then be assigned to the array in the order that they are dragged on to the array icon.
- 6) If you incorrectly selected the wrong disk for the array at this time, simply highlight that disk in Tree View, right-click the mouse, and select Delete from the pop-up menu. The disk will be removed from this array.
- 7) Repeat step 5 until the array is completed.
- 8) Press the “Commit” button and reboot the Operating System for a new array.

Deleting an Existing Array

NOTE: Deleting an array may result in the loss of all data contained on it. Be sure to back up any needed data before deletion.

- 1) Select the icon of the array  you wish to delete
- 2) Right-click the icon and choose "Delete" from the pop-up menu.
- 3) Click the "Yes" button if this is an array you wish to delete. Reboot.

Setting Up E-mail Alert Notification

- 1) Double-click on the Message Agent icon  from where you want to receive email alert notifications. This will reveal the Message Agent main screen.
- 2) Check the "Email alert on error" box in the Information View (if it is not already checked).
- 3) Type in the SMTP server address for your mail server in the "SMTP server" field. Your network administrator will have this information if you do not know it.
- 4) Click the "Change" button next to the input field to update SuperCheck configuration for the SMTP server.
- 5) Next, set the email address for this machine. The Information View is located below the "Email alert on error" pane. E-mail reports will seem to originate from this address. Recipients of these messages may reply to this address.
- 6) Click on the "Change" button next to the input field to update SuperCheck configuration for the Email ID.
- 7) Add users to the "Email Receiver List".

Adding an User to an Email Receiver List

- 1) Double-click on the Message Agent  . Add any new user(s) needed for email alert notification. This will open the Message Agent main screen. Find the Information View of the Message Agent screen.
- 2) If you have not already done so, set up email alert notification(s).
- 3) Under the "Email receiver list" segment of the Information View, type in the email address(es) of the user you wish to receive email alert notification(s) in the "Email address" field.
- 4) Click on the "Add" button to add any users to the list.

Removing an User From an Email Receiver List

- 1) Double-click on the Message Agent  from which you wish to remove a user for email alert notification(s). This will open the Message Agent main screen. Find the Information View of the Message Agent screen.
- 2) Under the “*Email receiver list*” segment, select the email address in the scrollbox which you wish to remove.
- 3) Click on the “Remove” button or press the Delete key to remove the address from notification.

Adding/Removing Events to/from Email Alert Notification

- 1) Double-click on the Message Agent icon  to change the alert email notification events you wish to modify. This will open the Message Agent main screen. Find the Information View of the Message Agent screen.
- 2) Select the event you wish to modify in the “Event” column.
- 3) Right-click on the “Email” column and select “yes” to have this event send out an email notification. Select “no” to remove the item as an alert event.
- 4) When you are finished making changes, click the “Change” button next to the event selection scrollbox to update the SuperCheck configuration for email alert notification.

NOTE: The “Report” column is used for event reporting to the utility. Events may be accessed by right-clicking the IOPx (controller) icon .

Using SuperCheck™

This section is an in-depth software functionality reference for the SuperCheck RAID Manager. It contains step-by-step instructions on activating events, modifying values, and executing all major tasks.

5.1 Main Window

Upon starting SuperCheck, the main window will be displayed. The figure below shows an example of the entire hierarchical tree.

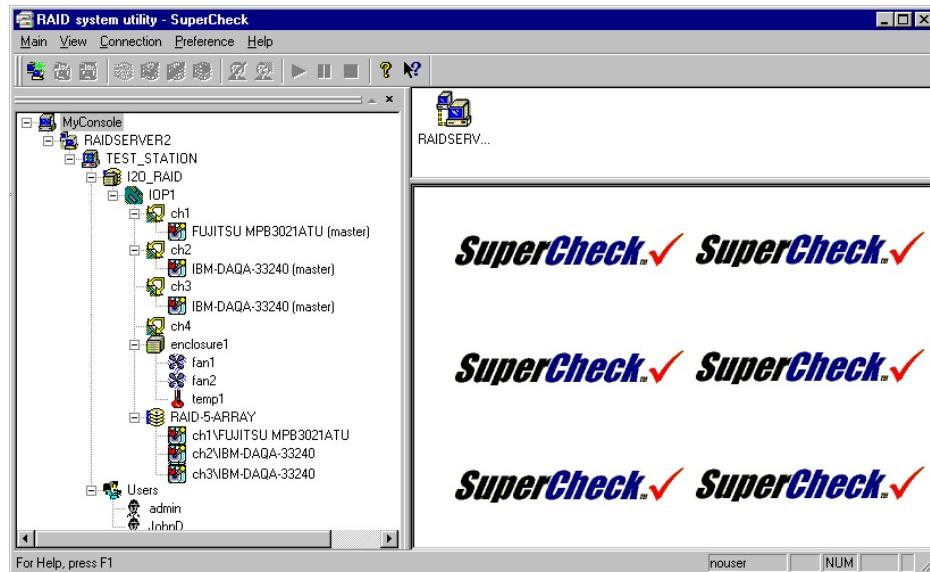


Figure 23: SuperCheck utility main window

Along the left side of the window is the **Tree View** which originates with “Console” and can display the entire SuperTrak66 system(s), down to individual hard drives within arrays and server-level components (see Chapter 4 for details on each level).

At the top right corner is the **Object View**, which displays icons representing lower-level devices that appear below the highlighted device in the Tree View.

The SuperCheck splash screen is displayed in the **Information View** of the window (the lower right hand section of the window.)

Using Tree View

The Tree View can display every element of your SuperTrak66 system. This menu behaves like the Tree View in Windows Explorer (Explorer shows logical drives and folders, etc...in a hierarchical menu structure). You may also choose to expand or collapse Tree View items, or hide/display the Tree View pane entirely.

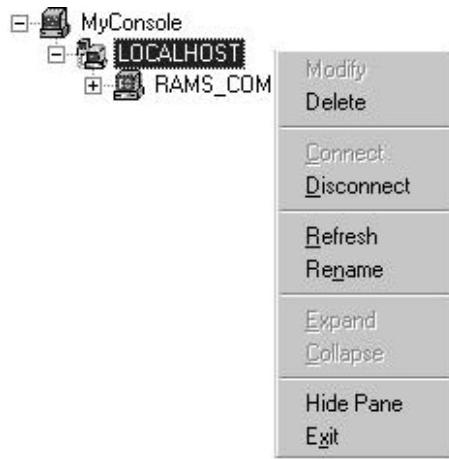


Figure 24: Using Tree View

Expanding/Collapsing Tree View Items

To open the items below a listed device, click once on the [+] to expand the tree, [-] to collapse it. Double-clicking an item will not only cause the tree to expand down, but will also add the sub-items as icons in the Object View (see Figure 24 above).

Hiding/Displaying Tree View

Users may hide or show the Tree View pane of the SuperTrak66 Main interface at any time. Navigation can be performed via the Object View and/or pull-down menus. The hide/show function is accessed either by right-clicking any icon in the Tree View and selecting "Hide" or "Show". The same function may be used with the pull-down menu with any icon highlighted in Tree View (see Figure 24 above).

Renaming Displayed Levels

Users may rename the array and server names at the Tree View depending on their administration rights.

The renaming function is accessed either by right-clicking the appropriate icon in the Tree View and selecting "Rename" or by using the pull-down menu with the appropriate icon highlighted in Tree View.

Console & Server level rights: no special user or administrative rights are required since the name applies only to the user's system.

Array & User level rights: requires administrative (password-protected) rights since other Agents can see the array and potentially access the levels.

Refreshing SuperTrak Monitoring

Users may manually refresh SuperTrak66's display of the server environment shown in Tree View at any time to immediately see system changes anywhere in the network. The default refresh rate is once every 10 seconds.

The refresh function is accessed either by right-clicking icons in the Tree View and selecting "Refresh" or by using the pull-down menu with the appropriate icon highlighted in Tree View. It can also be accessed using the F5 function key.

Using Object View

The device icons generated by a double-click in the Tree View will be displayed in this portion of the window. Double-clicking an icon in the Object View highlights the item corresponding to it in the Tree View and changes the Object View to include any items directly connected to the icon.



Figure 25: Using Object View

NOTE: Icons shown in this pane of the window which have a red "X" over them represent an error or failed device.

Using Information View

The lower right portion of the SuperCheck main window is the Information View. The Information View contains different items (text boxes, list boxes, information fields and buttons), depending on which device is highlighted in the Tree View.

System Information			
Name	Version	Date	Size
IxWorks 1.1 FCS	0201		0
PROMISE RAID ISM	b4pr	02/24/2000	283556
PROMISE IDE HDM	b4pr	02/24/2000	59579
Processor	Intel 80960 series		
Memory Size (MB)	8.00		
HardWare	MORY-8		

Figure 26: Using Information View

Using Status Bar



Figure 27: Using Status Bar

SuperCheck's status bar is no different than any other Windows program. The status bar is located just below the Main Window of SuperCheck and provides status information of various administration functions such as (Ready, Rebuilding, etc.). The Status Bar can be displayed or hidden using the View pull-down menu.

5.2 SuperCheck Pull-Down Menu/Toolbar

SuperCheck uses a context-sensitive pull-down menu and optional toolbar with icons located at the top of the SuperCheck Utility Main Menu. They provide access to all major SuperCheck management functions.

Using Pull-Down Menu Items

In addition to the options accessed from Tree View and the Information View, users may access the same functions from pull-down menu selections located at the top of the SuperCheck Utility Main Menu. The options available are context-sensitive depending on which icon in the Tree View is highlighted (see example below). Certain functions that are unavailable will appear "greyed out." Directions on implementing functions and options in this Manual are based on the Tree View unless otherwise indicated.



Figure 28: Array pull-down menu

Using View Pull-Down Menu



Figure 29: Using View Pull-Down Menu

By checking or unchecking items with the mouse, this pull-down menu selects or deselects the appearance of the Toolbar icons, Status Bar, and/or Tree View window.

Using Connection Pull-Down Menu

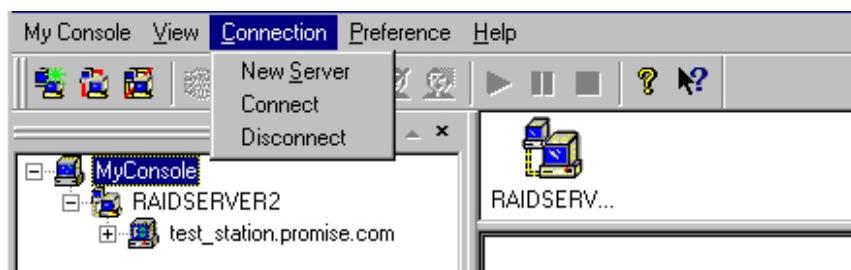


Figure 30: Using Connection Pull-Down Menu

The Connection pull-down screen is available at all times regardless of which icon is highlighted in the Tree View of the SuperCheck utility. From here, a user may select from the following functions:

New Server

Creates a new server (see section 5.3 Console Functions on page 59)

Connect

Establishes a connection with an already-created SuperTrak66 server (see section 5.3 Connecting The Message Server on page 61)

Disconnect

Temporarily halts the connection with a SuperTrak66 server (see section 5.3 Disconnecting A Message Server on page 61)

Using Preference Pull-Down Menu

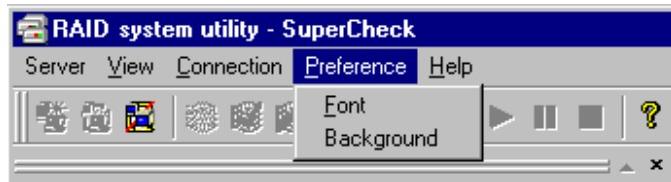


Figure 31: Using Preference Pull-Down Menu

The Preference pull-down menu allows users or administrators to change the way the SuperCheck Main Window is displayed:

Font

Selecting Font allows you to select a different font and/or font color to display text

Background

Allows changing the color of the SuperCheck View panes.

Using Help Pull-Down Menu



Figure 32: Using Help Pull-Down Menu

Activated from the pull-down menu, the Help Menu allows users to access an online help index of how-to topics and functions, provides the SuperCheck version/build identification, and can activate a demonstration of all Main Window screens.

About SuperCheck

Pull-down menu item provides information on the version and build number of the SuperCheck utility.

Viewing Autodemo

A level-by-level automatic display of each screen in the SuperCheck Main Window, starting at the Console level in Tree View. You may also control the Autodemo via the keyboard as follows:

Start/Stop AutoDemo	press F12 key
Next Item	press F10
Previous Item	press F11

Using Help Topics

Activated from the pull-down menu or toolbar icon, online help offers information on functions and how-to items accessible from the SuperCheck utility. There are two major areas, Contents and Find.

Contents

Lists SuperCheck functions either by Interface Component (i.e. Tree View level item or Information View Menus) or by how-to actions (Starting Up, Managing Arrays, etc.).

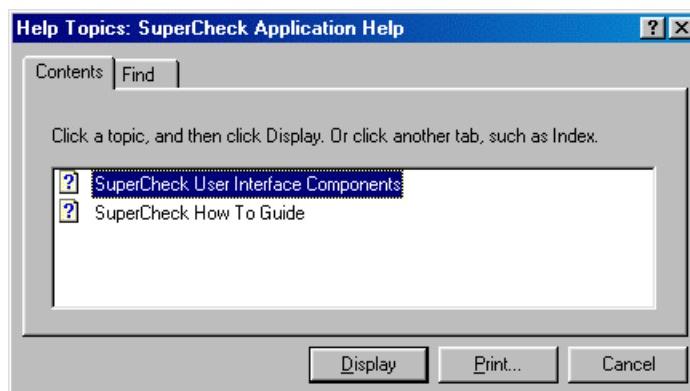


Figure 33: Using Help Contents Window

Find

Creates a comprehensive index based on each word and topic found in the Help file.

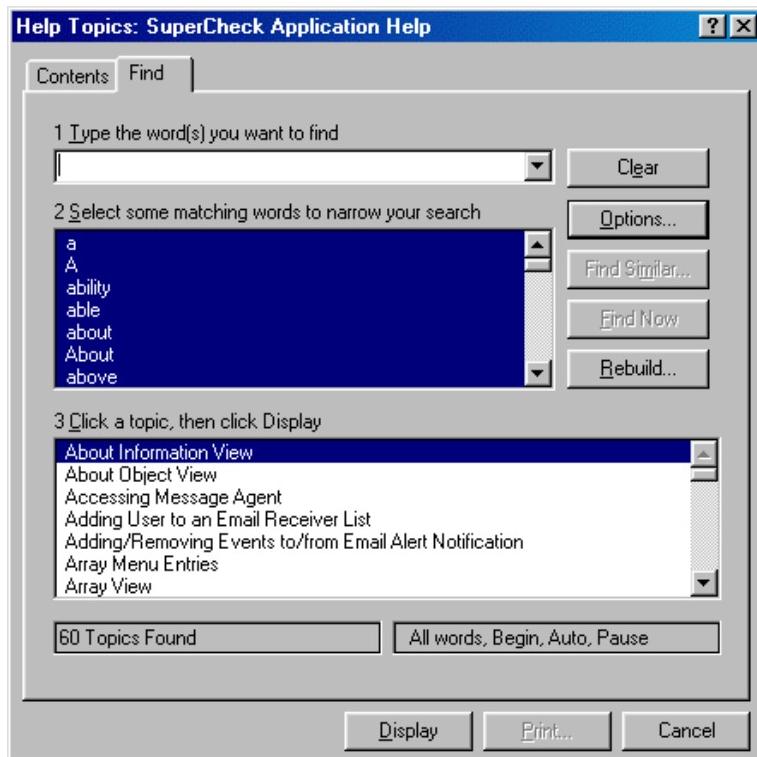


Figure 34: Using Help Find Window

Using SuperCheck Toolbar



Figure 35: Toolbar icons

Toolbar icons and their associated functions are available depending on which of the items in the Tree View is highlighted as described in the descriptions below.



New Server: available when the Console icon is active. It allows the user to create a New Server on the SuperTrak66 system. The user must have “User Account Rights” (see p. 59 for details).



Connect Server: available when a particular Server icon is active. It allows the user to recognize the Server on the SuperTrak66 system. The user must have “User Account Rights” (see p. 61 for details).



Disconnect Server: available when the Server icon is active. It allows the user to take the Server offline from the SuperTrak66 system. The user must have “User Account Rights”. Disconnect is the only proper way to shut down the SuperTrak66 server in order to perform physical maintenance. This also logs off the Server from being recognized by any other SuperCheck monitoring utility online (see p. 61 for details).



Rebuild Array: available when an Array icon for a RAID Level 1, 3, or 5 is active. This begins the process of restoring data to a replaced drive from remaining drives in the array. The user must have “User Account Rights” (see p. 91 for details).



Synchronize Array: available when an Array icon for a RAID 1 or 0+1 level array is active. This allows user to write data from primary mirrored drive(s) to secondary mirrored drive(s) to assure both have identical data as a maintenance procedure. The user must have “User Account Rights” (see p. 94 for details).



Delete Array: available when an Array icon is active. This allows the user to delete the array from the SuperTrak66 system. The user must have “User Account Rights” (see p. 90 for details).



Create Array: available when an Message Agent icon is active. This allows user to create a New Array on the SuperTrak66 system. The user must have “User Account Rights” (see p. 80 for details).



Delete User: available when a User icon is active. This allows an administrator to delete the user from SuperTrak66 monitoring/alert email rights access. Only administrators may delete an user (see p. 75 for details).



Create User: available when a Message Agent icon is active. This allows the admin to create and add a new user to the SuperTrak66 system for monitoring/alert email. The user must have “User Account Rights” (see p. 66 for details).



Start/Continue Button: available when Array icon is active. This allows start/continuation of rebuild (see p. 94 for details) or array synchronization process (see p. 95 for details).



Pause Button: available when Array icon is active. This allows pausing of rebuild (see p. 94 for details) or array synchronization process (see p. 95 for details).



Stop Button: available when Array icon is active. This allows start/continuation of rebuild (see p. 94 for details) or array synchronization process (see p. 95 for details).



About SuperCheck: provides information on the version and build number of the SuperCheck interface (see p. 54 for details).



Help Button: loads SuperCheck' s Help contents and index. (see p. 55 for details).

5.3 Console Functions



The SuperCheck "Console" in the SuperCheck utility main screen represents the system from which you are monitoring your SuperTrak66 system(s). Its icon appears at the top of the Tree View (see below). When this icon is highlighted, you may also access all Console functions from the pull-down menus at the top of the main menu or context-sensitive menus enabled by a right-click of the mouse.

Creating A New Message Server

- 1) Right-click on the "MyConsole" icon in the Tree View area of the SuperCheck main menu (or left-click "MyConsole" from the pull-down selections at the top of the main menu).
- 2) Left-click on "New" in the console menu, then choose "Server" in the pop-up menu.

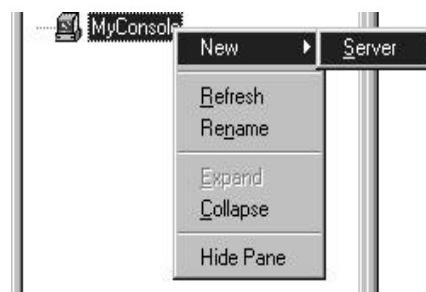


Figure 36: Creating New Message Server pull-downs

- 3) In the Message Server Information View (see below), enter "Label," "Server Name" and "IP address" information. If you are not sure of the actual Server name, try just entering the IP address of the machine where the Message Server is installed. Enter 127-0-0-1 for a local machine connection.
- 4) Click the "Confirm" button.

Label	RAIDSERVER1
ServerName	New Server
IP Address	127 . 0 . 0 . 1
Version Information

Confirm Reset

Figure 37: Labeling New Message Server

Viewing Console Object icons

Once a Message Server has been created, double-clicking on the Console icon will display the associated “Message Server” icon(s) in the Object View. These represent the message server(s) connected to the SuperTrak66 console (see figure below). Clicking on the server icon(s) will display the Information View screen for the Message Server (see above).



Figure 38: Console Object View

5.4 Message Server Functions



The “Message Server” icon in the SuperCheck Tree View represents the server that acts as a “gateway” into a particular SuperTrak Pro system. The system may be composed of one server (the “gateway” server itself), or multiple servers. Message Servers may or may not have a SuperTrak66 controller card and hard disk array installed. Once created, the Message Server will automatically recognize all associated Message Agents that exist on the network. Through the Message Server icon, users can delete a Message Server, and connect/disconnect the Message Server from system. When this icon is highlighted, you may also access all Message Server functions from the pull-down menus at the top of the main window.

Deleting A Message Server

NOTE: If the server's designated IP address is changed because of a physical move or system change, the Message Server must be deleted and a new Message Server created with the new IP address and Server name.

- 1) Right-click on the “Message Server” icon in the Tree View area of the SuperCheck main menu (or left-click "Server" from the pull-down selections at the top of the main menu with the "Server" icon highlighted in Tree View).
- 2) Left-click on “Delete” in the pull-down/pop-up menu.

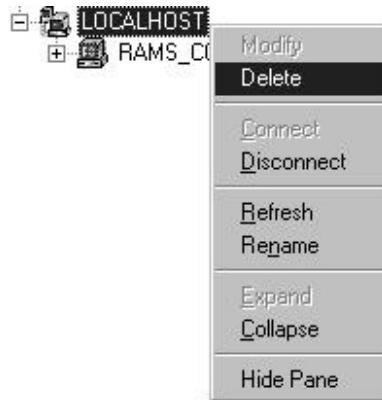


Figure 39: Deleting Message Server pull-down

- 3) Click the “Confirm” button. The Server name and icon will be removed from the SuperCheck Main menu screen. This change will take effect on the next system reboot.

Disconnecting A Message Server

- 1) Right-click on the “Message Server” icon in the Tree View area of the SuperCheck main menu (or left-click "Server" from the pull-down selections at the top of the main menu with the "Server" icon highlighted in Tree View).

- 2) Left-click on "Disconnect" in the pull-down/pop-up menu.

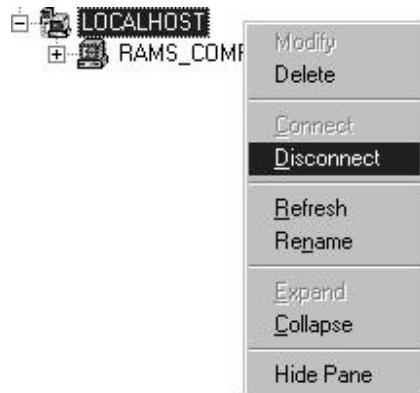


Figure 40: Disconnecting Message Server pull-down

- 3) Click the "Confirm" button. The Server connection from the SuperTrak66 will be taken "offline." This change will take effect immediately.

NOTE: *Disconnecting the Server from the SuperTrak66 system prevents all SuperCheck utilities from accessing data contained on the Server array(s). This allows server maintenance to be performed but does not require deleting and re-creating the Message Server data. This is also the only way to officially log out the SuperTrak66 Server from the system. The utility will still see the Message Server, but in a "disconnect" status.*

Connecting The Message Server

- 1) To connect a Message Server that has been taken "offline", right-click on the "Message Server" icon in the Tree View area of the SuperCheck main menu (see below) or left-click "Server" from the pull-down selections at the top of the main menu with the "Server" icon highlighted in Tree View.
- 2) Left-click on "Connect" in the pull-down/pop-up menu.

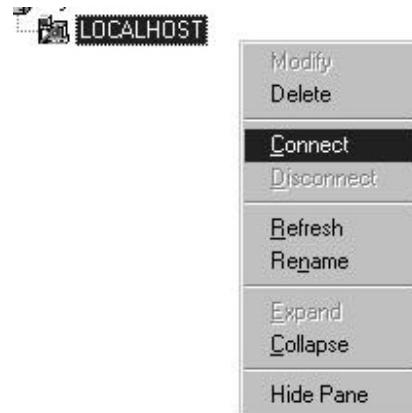


Figure 41: Connecting Message Server pull-down

Switching to Another Server

- 1) To connect to a Message Server in a different location, make sure the Message Server icon is highlighted, then enter the IP address and label name of the desired server in the Information View (see Figure 37).
- 2) Click the “Confirm” button. You will now be connected to another SuperTrak66 system location’s Message Server and perform administrative functions (assuming you have password authority).

NOTE: You would switch Message Servers if the physical location of a Server or IP address has changed.

Viewing Message Server Object icons

Double-clicking on the “Message Server” icon will display the SuperTrak66 system server icons connected to the highlighted message server within the Object View (see top of figure below). Clicking on the individual icons here reveals the Message Agent Information View (see section 5.5).

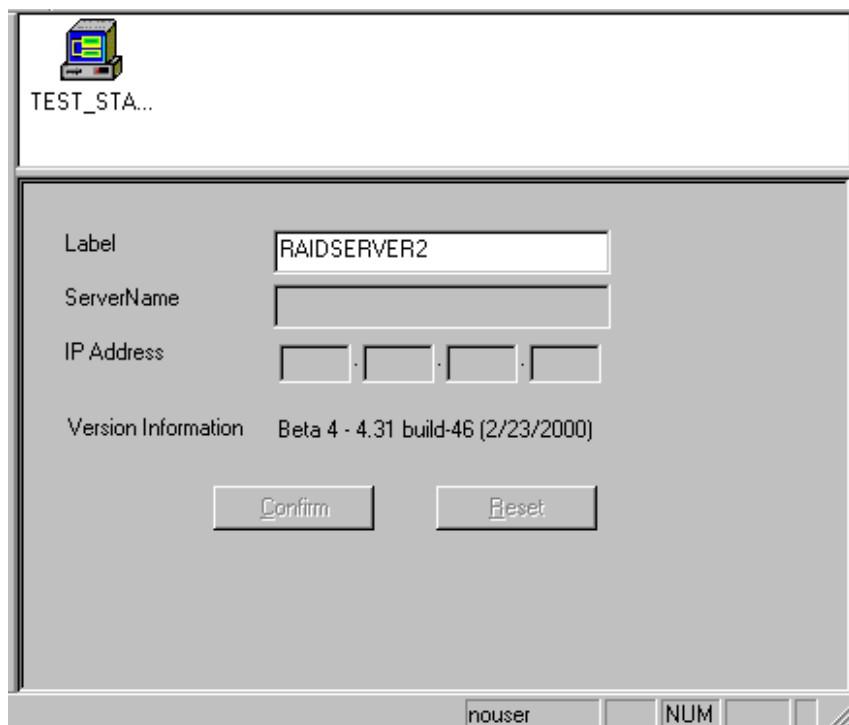


Figure 42: Message Server Information View

Viewing Message Server Information

In the Information View of the SuperCheck window, several text boxes will appear. The fields for the Message Server Information View are defined as follows:

Label

In the “Label” text box, the name of the message server can be entered or modified, and will be displayed within the Tree View.

Server Name

The “Server name” text box contains the internal operating system name of the message server currently highlighted. This entry is obtained from the server itself and cannot be modified.

IP Address

This series of four text boxes contains the IP address of the message server currently highlighted. The IP address field is greyed out so you can not change the IP address.

Version Information

This field contains information pertaining to the build version of SuperCheck currently operating. This field cannot be modified.

Confirm / Reset

These buttons relate to the text boxes titled "Label," "Server name," and "IP address". Selecting "Confirm" saves the entries in these fields after they have been modified. Selecting "Reset" clears the text boxes and allows the user to modify old information.

5.5 Message Agent Functions



The “Message Agent” icon(s) in the SuperCheck utility main screen represent the SuperTrak66 servers connected to the Message Server in a particular SuperTrak66 system. There may be only one member in a network (the server itself), or there may be additional servers connected (remote systems). The Message Agent allows you to create new users, establish events for notification, create e-mail addresses for notification, and schedule array synchronization.

Accessing Message Agent

- 1) By clicking on any Message Agent icon in Tree View or Message Server Object View (if you are not already accessing the Agent), you will be prompted for password entry.



Figure 43: Setting Password screen

- 2) Enter the account/password for this system to access the Agent and gain permission to the SuperTrak66 system.
- 3) Checking the “Save password” box will remember the password for this Message Agent. Future access to this Message Agent from this terminal will not require password input.
- 4) Click “OK” to submit the password and access the Message Agent.

NOTE: This Message Agent will remain accessible for the duration of the SuperCheck session. Note: the utility will NOT automatically log out. You must perform this manually.

Creating A New User

- 1) Log into the Message Agent for which you would like to add a New User.
- 2) Right-click the “Users” icon in the Tree View area of SuperCheck Main Menu. Then left-click on “New” and select “User” from the pop-up selection (see below).

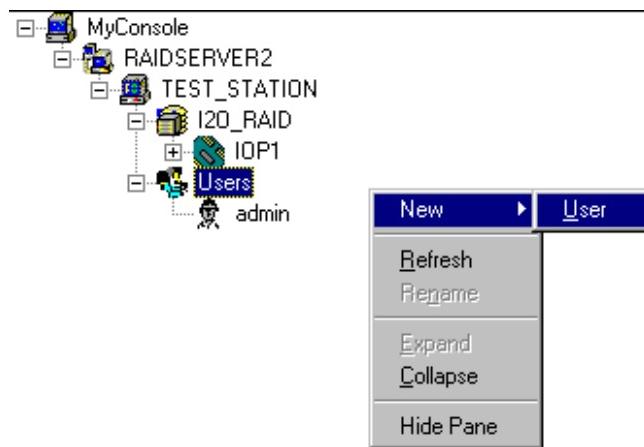


Figure 44: Creating New User screen

- 3) A “New User” icon will appear. Double-click on it to view the user Information View.
- 4) Modify all the user information according to the desired access level for the new user (see section 5.6 for more details on “User Rights”)
- 5) If the user has any level of administrative control, and the server is connected to a WAN, make sure to set a password.
- 6) Click the “Commit” button to update the user information.

Setting Up E-mail Alert Notification

- 1) Double-click on the Message Agent icon from where you want to receive email alert notifications. This will reveal the Message Agent main screen. Find the Information View of the Message Agent screen similar to below.



Figure 45: Setting E-mail Alert box

- 2) Check the "*Email alert on error*" box in the Information View (if it is not already checked).
- 3) Type in the SMTP server address for your mail server in the "SMTP server" field. Your network administrator will have this information if you do not know it.
- 4) Click the "Change" button next to the input field to update SuperCheck configuration for the SMTP server.
- 5) Next, set the email address for this machine. The email address section is located below the "Email alert on error" pane. It appears below. E-mail reports will seem to originate from this address. Recipients of these messages may reply to this address.

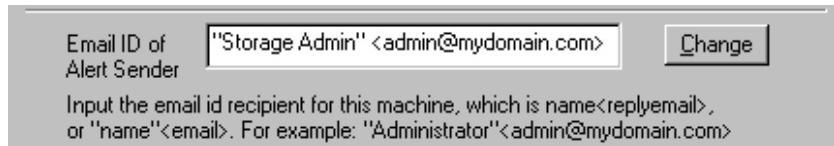


Figure 46: Setting Email ID window

- 6) Click on the "Change" button next to the input field to update SuperCheck configuration for the Email ID.
- 5) Add users to the "*Email receiver list*".

Adding an User to an Email Receiver List

- 1) Double-click on the Message Agent to which you wish to add a user for email alert notification. This will reveal the Message Agent main screen. Find the section of the Message Agent Information View screen similar to Figure 47.
- 2) If you have not done so already, set up email alert notification.



Figure 47: Setting Email Address window

- 3) Under the “*Email receiver list*” segment of the Information View, type in the email address of the user you wish to receive email alert notification in the “*Email address*” field.
- 4) Click on the “Add” button to add this user to the list.

Removing an User From an Email Receiver List

- 1) Double-click on the Message Agent from which you wish to remove a user for email alert notification. This will reveal the Message Agent main screen. Find the Information View of the Message Agent screen similar to figure below.



Figure 48: Removing User from E-mail alert

- 2) Under the “*Email receiver list*” segment, select the email address in the scrollbox which you wish to remove.
- 3) Click on the “Remove” button to remove the address from notification.

Adding/Removing Events to/from Email Alert Notification

- 1) Double-click on the Message Agent icon whose alert email notification events you wish to modify. This will reveal the Message Agent main screen. Find the Information View of the Message Agent screen similar to Figure 49.

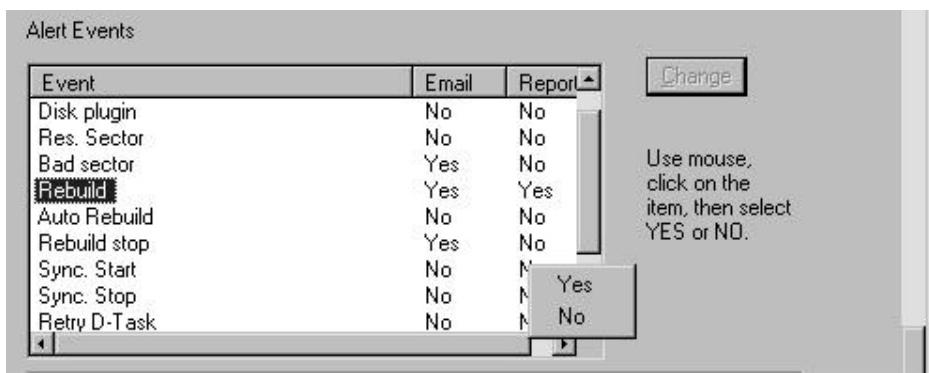


Figure 49: Adding/Removing Events for Email Alert

- 2) Select the event you wish to modify in the "Event" column.
- 3) Right-click on the "Email" column and select "yes" to have this event send alert email notification. Otherwise, select "no" to remove it as an alert event.
- 4) When you are finished making changes, click the "Change" button next to the event selection scrollbox to update the SuperCheck configuration for email alert notification.

NOTE: The "Report" column is used for event reporting to the utility. Events may be accessed by right-clicking the IPOx item.

Scheduling Array Synchronization

The Schedule Array Synchronization section (see below) allows the administrator to perform synchronization on a regularly-scheduled basis (by minute, by hour, by day, by week, or by month).

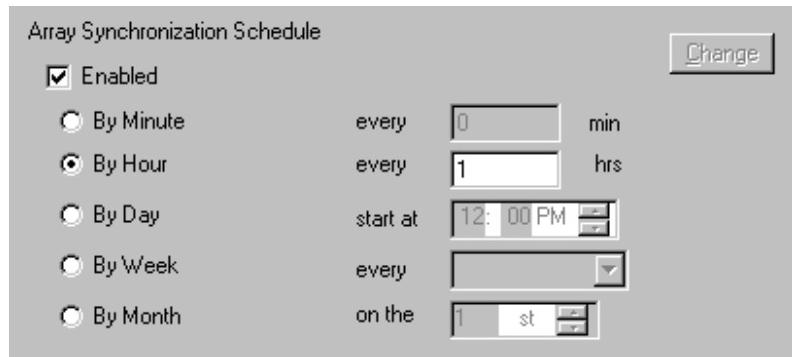


Figure 50: Scheduling Array Synchronization

- 1) To schedule synchronization, check the Enable box,
- 2) Select the scheduling options appropriate to your timetable
- 3) Click the "Change" button for the new schedule to take place.

NOTE: During Array Synchronization, users may continue to access the working array and perform normal PC functions. However, system performance will be slightly degraded and the process will take longer. Scheduling allows synchronization to take place at off-peak time periods.

Viewing Message Agent Object icons

Double-clicking on a "Message Agent" icon will generate "I_O_RAID" and "User" icons in the Object View. Left-clicking the I_O icon (left) will reveal specific Agent information in the Information View (see next section). Left-clicking User icon(s) will reveal User information.



Figure 51: Message Agent Console icons

Message Agent Information View

The Message Agent Information View (see Figure 52 below) is a very long screen and requires the use of the scroll bar to view all contained fields. Descriptions of the sections within the Message Agent Information View are as follows:

NOTE: To see all the fields in the Information View pane, minimize the Console View by placing the mouse cursor between the two panes so that the cursor changes to the window control cursor, hold down the left mouse button, and drag the boundary up until it disappears.

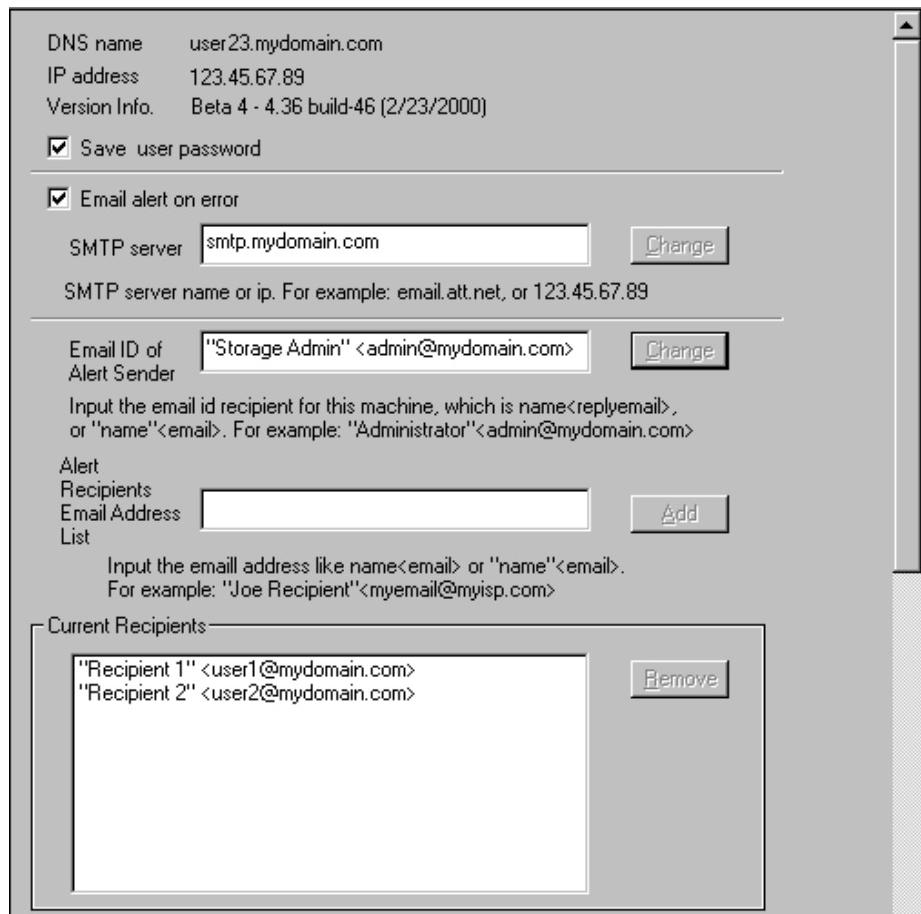


Figure 52: Message Agent Information View

DNS Name, IP Address & Version Info

These three fields correspond to the “Server name,” “IP address,” and “Version information” fields in the Message server information box above. They cannot be modified from here.

Save User Password

By checking this option, your password will be saved. The next time you use this console, you will not have to input your password again.

Email Alert on Error

By checking this option, notification through email can be implemented when an error occurs on the currently highlighted SuperTrak66 Message Agent. You will then need to enter the email address information of the

SuperTrak66 server in question in the two boxes directly below (“SMTP server” and “Email ID”.)

SMTP Server

Enter the SMTP server name or IP address of your mail server. Contact the System Administrator if you need this information. If you need to change information already entered, press the “Change” button next to the text box after typing in any new information.

Email ID

Enter the email ID of the email location for the specific machine housing SuperTrak66 in this field. Email alerts will originate from this address. If you need to change information already entered, press the “Change” button next to the box after typing in any new information.

Alert Recipients Email Address List

Enter the email addresses of all individuals that need notification when an error occurs in a particular SuperTrak66 station in this field. If you are monitoring multiple SuperTrak66 station, you must create a separate email list for each SuperTrak66 station. After entering an email address, press the “Add” button to confirm the entry.

Current Recipients List

This area of the Information View contains all email addresses to which an alert will be sent in the event of an error occurring on the SuperTrak system. To add a new address, enter it in the box; to remove an address, highlight it and then press the “Remove” button.

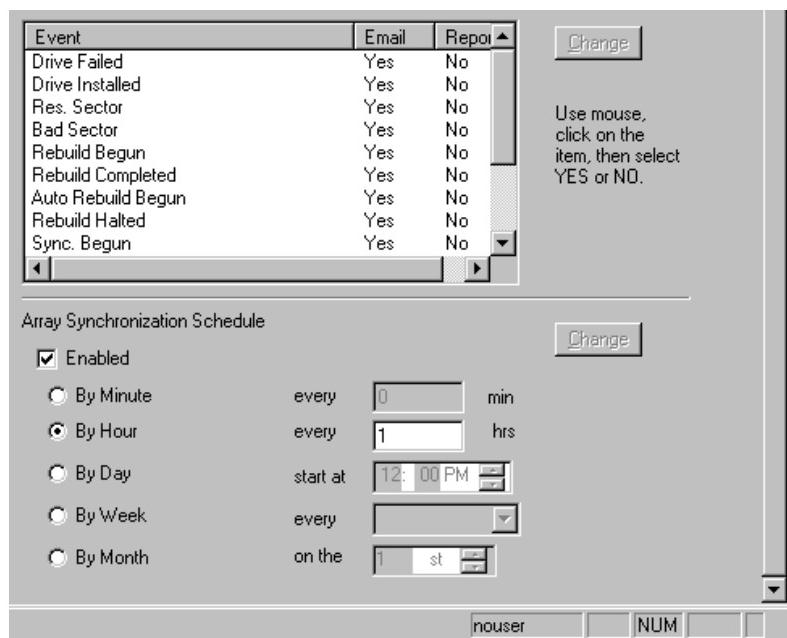


Figure 53: Alert Events & Schedule Synchronization screen

Alert Events

This area (see above) allows you to choose which types of events will generate an email alert and/or report. To do this, highlight a particular event in the "Event" column, then click once in the same row on either the "Email" or "Report" columns to toggle their active states between "Yes" and "No". The email notification sends email to the alert recipients. The "Report" column affects which events are stored in the Event List on the Message Agent. The Event list may be retrieved by right-clicking the IOPx item. It may also be deleted from here.

Schedule Array Synchronization

Array synchronization is a periodic maintenance procedure for mirroring (RAID 1, 0+1, 3, and 5) arrays to maintain data consistency on all mirrored drives. In operation, array synchronization compares data on the mirrored drives for any differences. If there are differences detected, data from the primary drive(s) is automatically copied to the secondary drive(s). This assures that all mirrored drives will contain the exact information.

5.6 User Management Functions



The “Users” icon in the Tree View is a folder containing a list of users and administrators who are permitted access to a particular SuperTrak66 system’s status.

Creating A New User

- 1) Right-click the “Users” icon in the Tree View area of SuperCheck Main Menu. Then left-click on “New” and select “User” from the pop-up selection.

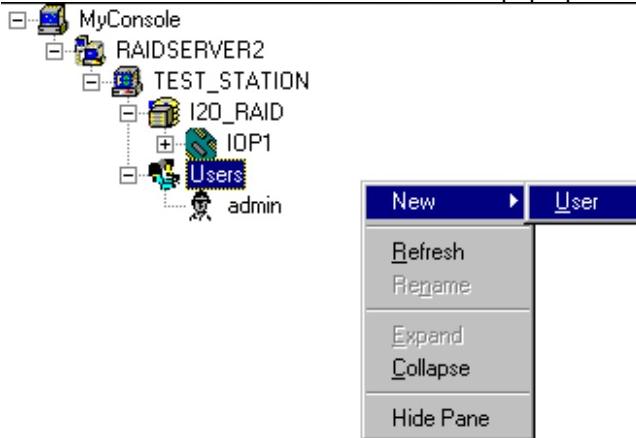


Figure 54: Creating New User pull-down

- 2) A “New User” icon will appear. Double-click on it to view the user Information View.
- 3) Modify all the user information according to the desired access level for the new user.
- 4) If the user has any level of administrative control, and the server is connected to a WAN, make sure to set a password.
- 5) Click the “Commit” button to update the user information.

Deleting an Existing User

- 1) Select the User icon in the Object View or Tree View representing the particular user you wish to delete.

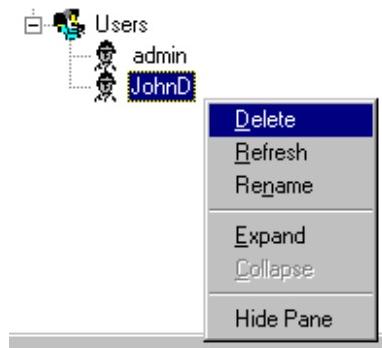


Figure 55: Deleting User menu

- 2) Right-click the icon and select “Delete” from the pop-up menu.
- 3) Choose “Yes” if this is a user you wish to delete.

NOTE: SuperCheck will not allow deletion of the last account with admin privileges. This protects the admin from being locked out of the system and having to re-install SuperCheck.

Changing Passwords

- 1) Double-click the Admin/User icon in Tree View or Object View of the person for which you want to change passwords.
- 2) In the Information View (see), type the new password in the labeled box and confirm the password in the box below it.
- 3) Click “Commit” to enable the change or continue with other selections.

Assigning User Rights

- 1) Double-click the Admin/User icon in Tree View or Object View of the person for which you want to assign Array Rights.
- 2) In the Information View (see figure 5-23), you may select from the following rights: "Create Array," "Delete Array," "Maintain Array," "Array Statistics," "Adapter Maintenance," and "User Account." Click the appropriate boxes to enable/disable these rights. See section 5.6.5 for a more detailed explanation of these rights.
- 3) Click “Commit” to enable the change.

Viewing User/Admin Object icons

Double-clicking on the “Users” icon will generate “User” and “Admin” icons in the Object View (see figure below). Clicking on the individual “User” and “Admin” icons will reveal their access rights (see page 77).



Figure 56: User icons in Object View



The “Admin” icon in the Tree View represents the configuration of a user who has “User Account Rights” on a particular SuperTrak66 system. Double-clicking on the icon will display an “Administrator Information View”, as shown in Figure 57 next page.

NOTE: No Object View will appear since this is the lowest level in this portion of the Tree View.

User Information View

The User Information View screen (see Figure 57) offers the admin/user the ability to assign passwords, array maintenance, adapter maintenance, and user account rights. What follows is a description of the fields in the Information View.

User Name

This text box contains the name of the administrator / user.

Password

The “Password” text box contains the (mask hidden) password of the administrator / user.

Confirm Password

The “Confirm password” text box is used when changing the password. After entering the new password in the “Password” text box, enter it again in the “Confirm password” text box to ensure accuracy.

Array Administration Rights

By checking the boxes, the user will be able to create, delete, and maintain arrays, and view array statistics.

NOTE: Array maintenance, creation, and deletion must be performed locally on the actual machine where the SuperTrak66 card is physically located. Such functions are greyed out on a remote monitoring station.

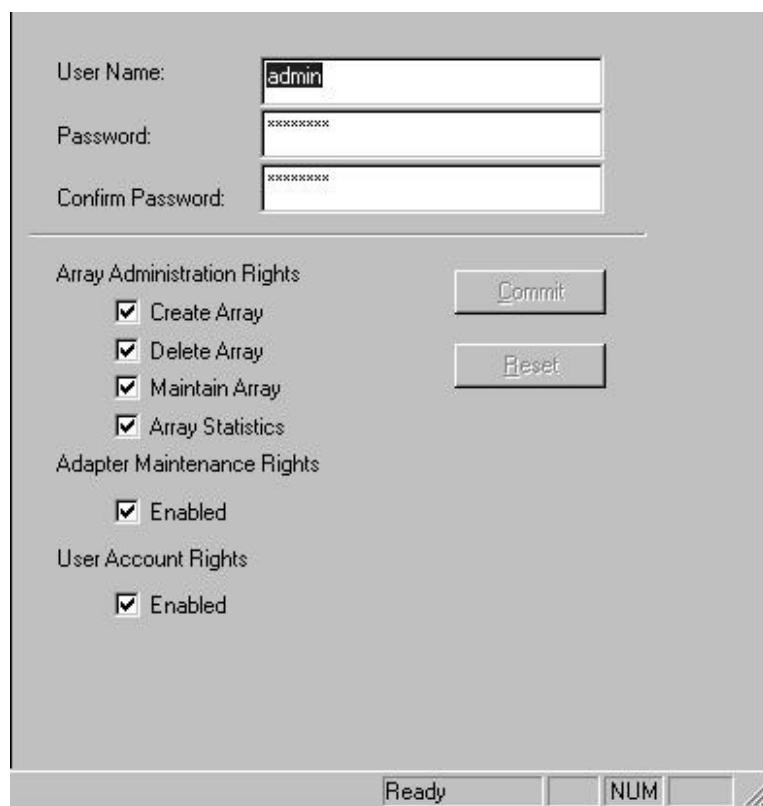


Figure 57: User Information View

Adapter Maintenance Rights

By checking the “Enabled” box, the user will be able to modify cache and performance tuning parameters for the SuperTrak66 card.

User Account Rights

By checking the “Enabled” box, the user will be able to assign or modify user “rights” (such as “Array Administration” or “Adapter Maintenance” rights) of the users listed. This user will also be able to add or remove other users.

Commit / Reset

If, as an administrator / user, you need to change your password, change it as described above. If a mistake is made, press the “Reset” button to revert to previously saved values. After correctly typing the new password text, press the “Commit” button. The new password will then be operational.

5.7 I₂O RAID Functions



The “I₂O RAID” icon(s) in the Tree View represents the I₂O component installed on a Message Agent. Through here, you can identify model information and hardware/firmware versions.

Viewing I₂O RAID Object icons

Double-clicking on an “I₂O RAID” icon generates the “IOPx” icon in the Object View (see below).



Figure 58: I₂O RAID Object View

Information View

In the Information View (see below), two fields are displayed: “Type” and “Version information” which reflect the model number, name of the controller installed in the system, and version information concerning the hardware and firmware. No fields may be modified here.

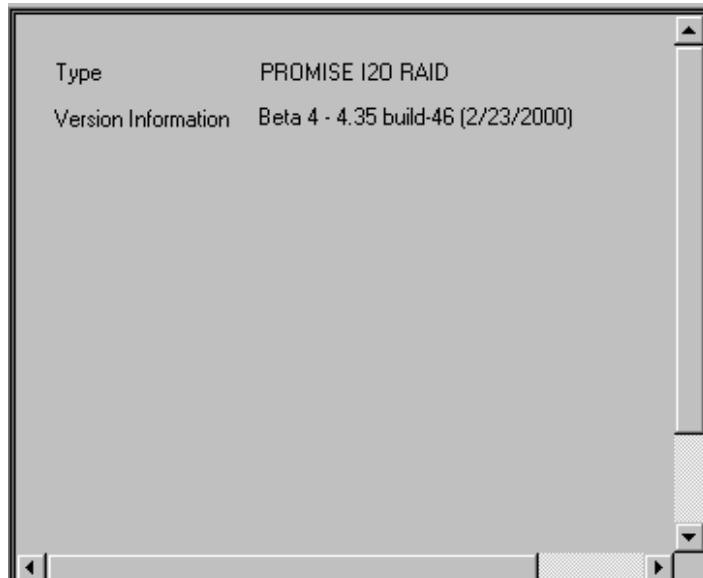


Figure 59: I₂O RAID Information View

5.8 IOPx Functions



The “IOPx” (I/O Platform #x) icon(s) represent individual SuperTrak66 controller cards which may be installed as part of a single or multiple SuperTrak66 server system as seen by SuperCheck. From here, users may create a new array, can view information on specific servers and SuperTrak66 cards. In addition, the user may choose to alter how often and what conditions of all SuperTrak66 cards' onboard memory will flush their data during operation

Creating a New Array

- 1) Select the IOPx icon in the Tree View where you want to create an array.
- 2) Right-click the icon. Click “New”, then select “Array” from the pop-up menu as show in figure below.

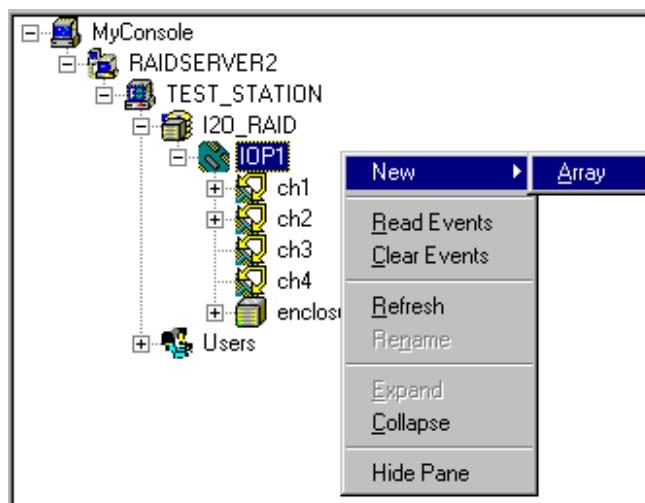


Figure 60: Creating New Array pull-down

- 2) Enter the pertinent information in the “Name,” RAID Level” and “Block Size” fields (see figure below)

Name	<input type="text" value="MY_RAID5_ARRAY"/>	<input type="button" value="Commit"/>
RAID level	<input type="button" value="RAID5(PARITY-DISTRIBUTED)"/>	
Block Size (KB)	<input type="text" value="64"/>	<input type="button" value="Reset"/>
Array Size	0 MB	
Array Status	functional	

Figure 61: Labeling New Array window

- 4) The new array icon will be created in the Tree View, titled with the text entered into the “Name” field.
- 5) Drag and drop any unassigned drive icon within the Tree View area on top of the array icon which you have just created in order to add the drive to the array (unassigned drive icons do not have a red arrow in the upper left corner). Hard drive(s) will then be assigned to the array in the order that they are dragged on to the array icon.
- 6) If you have incorrectly selected a disk for this array, highlight that disk in Tree View, right-click mouse, and select Delete. The disk will be removed from this array.
- 7) Repeat step 5 until the array is completed.
- 8) Press the “Commit” button and reboot the Operating System to establish the new array.

Reading/Clearing Events from SuperTrak Memory Buffer

SuperCheck allows the user to read alarm/maintenance events maintained by SuperTrak66's onboard memory log, print a listing of these events, save the log to a file, or clear all events from memory. Clearing events should be performed on a periodic basis in order to reduce the length and size of the event buffer log.

- 1) To read events from the log, make sure the IOPx icon of the correct SuperTrak66 controller card in Tree View is highlighted.
- 2) Left-click "Read Events" or "Clear Event" or use pull-down menu at top.
- 3) Clicking "Read Events" will enable the Event Viewer screen (see figure below)

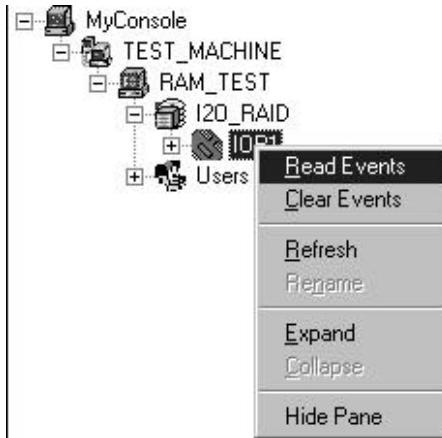


Figure 62: Read/Clear Events pull-down

Using The Event Viewer

After "Read Events" has been selected, the Event Viewer window appears (see below). The Event Viewer lists events according to type, date, status, and array status. The entire events log is not shown or accessed from the memory buffer since it can be quite large.

- 1) To sort the list according to category, click on the heading names
- 2) To navigate through the events listing, use the "Next 10" buttons to go forward through events.
- 3) To capture all events contained in the buffer prior to printing, click on the "Get All" button.
- 4) To save the Events log to a file, click the "Save As" button. Input a filename and click "Close" button.
- 5) Press "Finish" button to close the Event Viewer and return to the SuperCheck Main Menu.

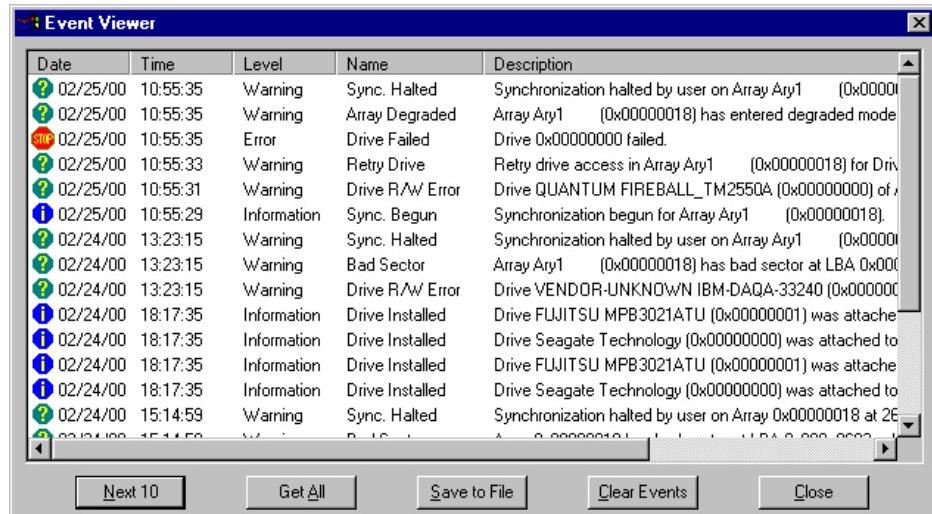


Figure 63: Events Viewer screen

Changing Flush Memory Timing

As part of the System Cache Policy settings, a user is able to change the amount of time the SuperTrak66 card holds cached data in its onboard memory.

- 1) Double-click the IOPx icon in Tree View.
- 2) In the Information View (see Figure 65 on p. 84), enter a new number to change the maximum number of seconds in which data in SuperTrak66's cache memory is held before being cleared.
- 3) Click "Commit" to enable the change or continue with other selections.

Changing Dirty Threshold Flush Start Setting

As part of the System Cache Policy settings, a user is able to change the maximum percentage of the total onboard SuperTrak66 memory that can be occupied by data that has not yet been written to disk (i.e. "dirty" data) before it is flushed from memory.

- 1) Double-click the IOPx icon in Tree View.
- 2) In the Information View (see Figure 65 on p. 84), enter a new number to change the maximum percentage of cache memory that can be occupied by "dirty" data before being cleared automatically.
- 3) Click "Commit" to enable the change or continue with other selections.

Changing Dirty Threshold Flush Stop Setting

As part of the System Cache Policy settings, a user is able to set the minimum percentage of memory allowed by SuperTrak66 for data that has not yet been written to disk (i.e. "dirty" data). This will stop SuperTrak66 from cleansing "dirty" data from memory beyond this percentage point.

- 1) Double-click the IOPx icon in Tree View.
- 2) In the Information View (see Figure 65 on p. 84), enter a new number to change the minimum percentage level of cache memory that can be occupied by "dirty" data and stop flushing automatically.
- 3) Click "Commit" to enable the change or continue with other selections.

Viewing IOPx Object icons

Double-clicking on this icon generates the components connected to the controller card – (from left) hard drive channels (first, second, and third icons below), enclosure components (fifth icon below), and arrays (sixth icon below) in Object View. Clicking on these icons will generate different information screens. See section 5.10 for drive channels, section 5.11 for enclosures, and section 5.12 for array.



Figure 64: IOPx Object View icons

IOPx Information View

The Information View displays the "IOPx Information View" as seen in Figure 65. The fields displayed in the "IOPx Information View" are defined as follows:

System Information

The "System information" section contains a list box which holds information concerning the controller's name, version number, timestamps, and size of file. Below it are three text fields: "Processor," "Memory size" and "Hardware". They show the type of CPU running on the controller, the amount of EDO memory present on the controller, and the type of interface chip on the controller respectively. Fields may not be modified.

The screenshot shows a software interface with two main sections: "System Information" and "System Cache Policy".

System Information

Name	Version	Date	Size
IxWorks 1.1 FCS	0201		0
PROMISE RAID ISM	b4pr	02/24/2000	283556
PROMISE IDE HDM	b4pr	02/24/2000	59579

Processor: Intel 80960 series
Memory Size (MB): 8.00
HardWare: MORY-8

System Cache Policy

Cache block size (KB)	16.00	<input type="button" value="Commit"/>
Total cache size (MB)	3.88	<input type="button" value="Reset"/>
R/W cache size (MB)	3.81	
Flush frequency timer (sec)	<input type="text" value="2"/>	<input type="button" value=""/>
Dirty threshold flush start (%)	<input type="text" value="80"/>	<input type="button" value=""/>
Dirty threshold flush stop (%)	<input type="text" value="5"/>	<input type="button" value=""/>

Figure 65: IOPx Information View

System Cache Policy

The “System cache policy” section contains a series of fields and text boxes as follows:

Cache block size (KB)

This field reflects the size of a unit block in the cache module. It cannot be modified from here.

Total cache size (MB)

This field indicates how much of the SuperTrak66 EDO memory is being used for R/W caching plus RAID XOR data space.

NOTE: Not all installed memory is available. A portion of cache memory is used to house I_O and other components of the Promise driver.

RW cache size (MB)

This field indicates how much of the SuperTrak66 EDO memory is being used for R/W caching only.

Flush frequency timer (sec)

This field shows how long (in seconds) a block of “dirty” data can remain in memory without being flushed. This field affects the cache configurations of any array attached to this IOP.

Dirty threshold flush start

This field indicates the flush start threshold percentage of dirty blocks. When the percentage of dirty blocks residing in memory exceeds the threshold, flushing will begin automatically. This field affects the cache configurations of any array attached to this IOP.

Dirty threshold flush stop

This field shows the flush stop threshold percentage of dirty blocks. When the percentage of dirty blocks residing in memory falls below the threshold, flushing will stop automatically. This field affects the cache configurations of any array attached to this IOP.

Commit / Reset

If changes are made to any of the “System Cache Policy” fields, press the “Commit” button to save the new settings, or “Reset” to make the values revert to their saved settings.

NOTE: The new settings will NOT be enabled until the next system boot.

5.9 Channel (chx) Functions



The "chx" icons represent particular drive channels (i.e. 1,2,3,4) on the SuperTrak66 controller card.

Viewing Channel Object icons

Double-clicking on the icon displays the icons of the hard drives connected to the channel in the Object View. There is no Information View for Channel (chx). Double-click on the drive icons in the Object View to obtain a drive's Information View (see p. 87).



Figure 66: Channelx's Object View



Assigned drive icon indicates drive that is already part of an existing SuperTrak66 array.



Unassigned drive icon indicates drive is not a part of an existing SuperTrak66 array. User may choose to leave drive to be used as a "hot" spare for a RAID 1, 3, or 5 level array. Unassigned drives may be assigned to create new arrays.

5.10 Hard Drive Functions



The hard drive icons represent hard drives connected to the SuperTrak66 controller card. The icons will be identified by the drive's make, model, status, size, DMA mode, and array configuration.

Hard Drive Information View

Since the hard drives are the last level in the SuperCheck Tree View, the Object View will not contain any information. Double-clicking on a hard drive icon updates the Information View with information regarding the particular drive, as depicted below.

Disk Model Name	SEAGATE ST313032A
Disk Status	OK
Disk Size	12.13 GB
Mode Setting	UDMA-2
Configuration	Assigned to an array

Figure 67: Hard drive Information View

This Information View simply displays information about the drive. None of the fields may be modified. The field descriptions are as follows:

Disk Model Name

This field shows the make and model of the disk.

Disk Status

This field shows the disk's operational status.

Disk Size

This field reflects the disk's total physical capacity.

Mode Setting

This field identifies what transfer mode the disk is set for read/writes

Configuration

This field shows whether the drive is free ("hot spare") or is assigned to an array.

5.11 Enclosure Functions

 The “EnclosureX” icon(s) represent the server-level monitoring capabilities of SuperCheck. From this icon, users may choose to enable alerting alarms during rebuilding and synchronizing of arrays. In addition, they may monitor enclosure status if the user has mounted the hard drives inside optional FastSwap66 or SuperSwap66 “hot” swap boxes (one drive per box) manufactured by Promise. FastSwap66 permits monitoring fan and housing temperature status. SuperSwap66 allows monitoring fan, temperature, and power.

Enabling Rebuild/Synchronization Beeper

System administrators may wish to enable or disable an audible beeper alert that the SuperTrak66 emits during rebuild of drives and synchronization of drives. The beeper warns the admin that the system is in a degraded status during these processes and alerts admin when the process is complete.

- 1) Double-click the Enclosure icon in Tree View.
- 2) In the Information View (see Figure 69), check or uncheck the Rebuild/synchronization Beeper box.

Viewing Enclosure Object icons

Double-clicking on the “Enclosure” icon in Tree View will display “fan” and “temp” icons in the Object View associated with that enclosure. Clicking on any of the Object View icons will display an Information View nearly identical to the “Enclosure” Information View. The only difference is that the Object View will not appear.

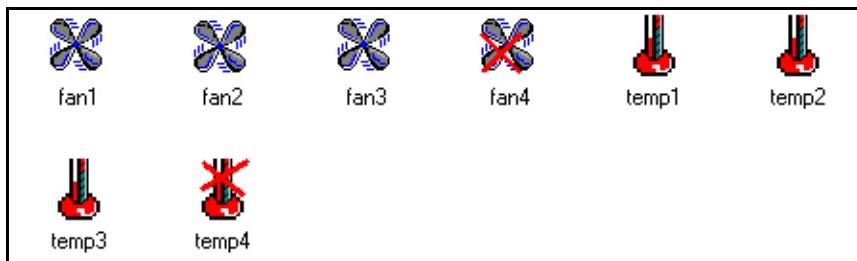


Figure 68: Enclosure Object View

Enclosure Information View

The Information View shows a graphical representation of the FastSwap66 or SuperSwap66 "hot" swap enclosures, along with the monitored components (fan, temperature, or power), as shown below.

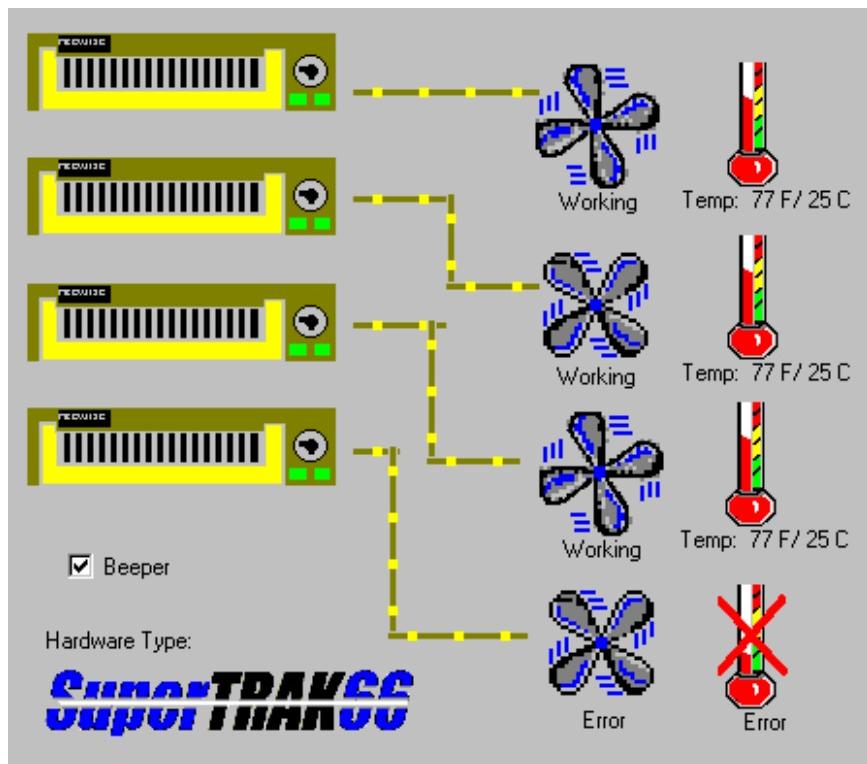


Figure 69: Enclosure Information View

The "fan" icons will display either a "working" fan, or a failed ("error") fan. The temperature icon will display a "working" thermometer (with the internal temperature in °F, or an "error" thermometer (if the temperature inside the enclosure exceeds 122° F). The "power" icons (not shown here) will display either "on" or "off" when using the SuperSwap66 enclosure.

5.12 Array Functions



The “Array” icon in the Tree View represents the architecture of a particular array. From here, admin users can delete an existing array, rebuild data to a replacement disk, synchronize data on mirrored drives, rename the array, turn on/off read and write back cache, and view cache statistics.

Deleting an Existing Array

NOTE: Deleting an array may result in the loss of all data contained on it. Be sure to back up any needed data before deletion.

- 1) Click the Array icon of the array you wish to delete in Tree View (see figure below)

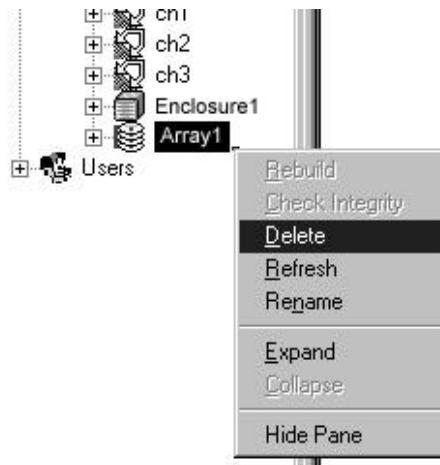


Figure 70: Deleting Array pop-up menu

- 2) Right-click the icon and choose “Delete” from the pop-up menu.
- 3) Click the “Yes” button if this is an array you wish to delete.
- 4) Reboot the system.

Rebuilding An Array

You will want to rebuild an array whenever a drive has failed and been removed from a RAID 1, 0+1, 3 or 5 array. For continuous operation, a replacement drive can be "hot" swapped while the attached system is operational.

NOTE: *In most cases, the rebuild process is initiated automatically -- either when a "hot" spare drive is available to the array or when a new hard drive is assigned to the array. The Rebuild Wizard screens (see page 92) will appear at this point.*

When rebuilding manually, data from the remaining working drive(s) can be reconstructed on to the target drive until data is restored. During rebuild, users may continue to access the working array and perform normal PC functions. However, system performance will be slightly degraded and the rebuild process will take longer.

NOTE: *Remember that during an array rebuild, the array does NOT provide data redundancy. It is recommended that all rebuilds be carried to completion.*

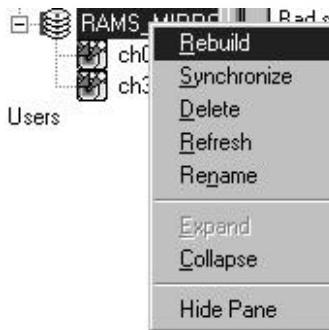


Figure 71: Array Pull-Down menu

- 1) Identify the location of the target (failed) drive within the SuperTrak66 housing based on the SuperCheck GUI. It will show which array is "critical" and identify which drive has failed (see Chapter 7 Troubleshooting for details).
- 2) Physically remove the failed drive and replace with an identical drive or one with the same capacity or larger.
- 3) Click the Array icon of the array where the drive is attached and where you wish to perform rebuild within the Tree View.
- 4) Right-click the icon and choose "Rebuild" from the pop-up menu. This will open the Rebuild Wizard Step 1 screen (see above).

Using Rebuild Wizard

Step 1

- 1) Select the correct Target drive which will receive data (see below) Make sure you select the blank new or replacement drive. The unselected drive(s) will contain "good" data. This will be the remaining working drive(s) of an array, or a system drive containing existing data that you wish to mirror.

NOTE: Make absolutely sure and double-check which drive is which. If data exists on the target drive, it will be over-written.

- 2) Click the Next button to proceed to Step 2 or Cancel button to stop.



Figure 72: Rebuild Wizard Step 1

Step 2

- 1) Confirm the Target or "Rebuild" disk by drive identification.
- 2) Click "Finish" button to initiate physical Rebuild, click the Back button to review Step 1, or Cancel button to Stop (see figure below)

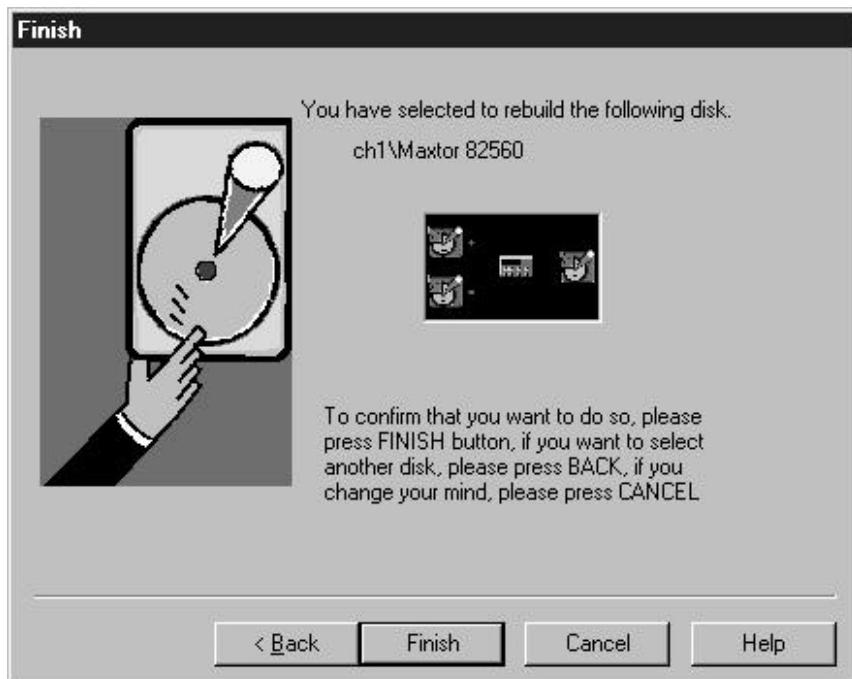


Figure 73: Rebuild Wizard Step 2

- 3) Once Array Rebuild has begun, you will be returned to the SuperCheck window. A progress window will appear below the Array icons in the Object View (see below). A similar progress bar will also appear in TreeView.

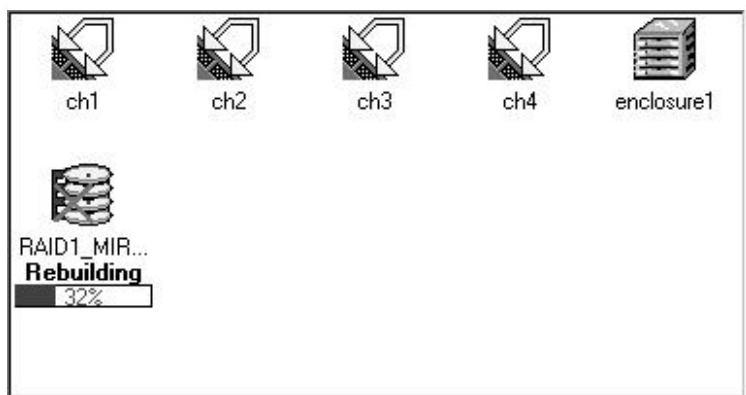


Figure 74: Array Rebuild Progress Bar

To Stop Rebuild

- 1) To halt the Rebuild process, right-click the Array # again.
- 2) The pull-down menu will appear showing the Stop Rebuild option as shown below (see below).
- 3) Once "Stop Rebuild" is selected, you will be asked to verify "Cancel Rebuild?"
- 4) Click "OK" to Stop or "No" to continue Rebuild.

NOTE: You may also start/continue, pause, or stop the process using the Toolbar icons.

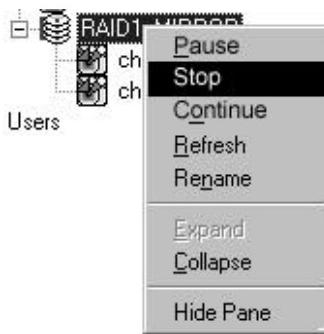


Figure 75: Stop Rebuild pull-down

Array Synchronization

Synchronization is a periodic maintenance procedure for Mirroring (RAID 1, 0+1, 3, and 5) arrays to maintain data consistency on all mirrored drives. In operation, array synchronization compares data on the mirrored drives for any differences. If there are differences detected, data from the primary drive(s) is automatically copied to the secondary drive(s). This assures that all mirrored drives will contain the exact information.

NOTE: Users may choose to schedule synchronization during unattended, off-peak hours (see page 74).

- 1) To manually synchronize drives from SuperCheck, choose the icon of the mirrored Array in Tree View.
- 2) Right-click the array icon, and choose "Synchronize" from the pull-down menu (see Figure 76) or select "Array" from pull down menu at top and choose "Synchronize" from the pull-down menu. This will initiate the process.
- 3) You will be returned to the main SuperCheck screen. A progress bar will display at the bottom of the window.

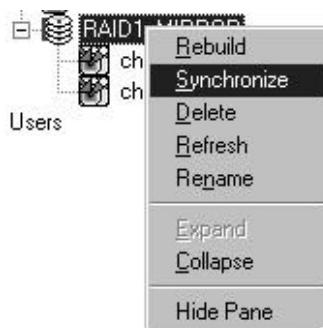


Figure 76: Array Synchronization pull-down

NOTE: During Array Synchronization, users may continue to access the working array and perform normal PC functions. However, system performance will be slightly degraded and the process will take longer.

Halting Synchronization

- 1) To halt the drive synchronization process, right-click the Array # again. The pull-down menu will appear showing the Stop Synchronization option as shown below. You may also temporarily pause the synchronization process and then continue.
- 2) Once "Stop" is selected, you will be asked to verify "Cancel Synchronization?"
- 3) Click "OK" to Stop or "No" to continue Synchronization.

NOTE: You may also start/continue, pause, or stop the process using the Toolbar icons.

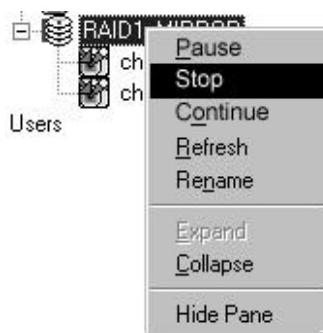


Figure 77: Halt Synchronization pull-down

Turning On Read Cache

Depending on the application or server operation in which the SuperTrak66 is used, admin/users may elect to turn on the read cache feature of the SuperTrak66 controller for performance purposes or disable caching entirely. Presence or absence of read cache may have a dramatic effect. Consult Chapter 6 RAID Management for more details.

- 1) To turn on read cache from SuperCheck, double-click the Array icon in Tree View.
- 2) In the Information View, left-click the arrow next to the Read Cache selection box and choose "On" (see p. 98).

Turning Off Caching

- 1) To turn off all caching from SuperCheck, double-click the Array icon in Tree View.
- 2) In the Information View, left-click the arrow next to the Read Cache selection box and choose "Cache Off". This will disable all caching functions from SuperTrak66 (see p. 98).

Turning On/Off Write Back Cache

Write-back cache affects system performance depending on server and application type. However, write-back cache data will be lost in the event of a system power outage or other event where the data has not yet been saved to disk. Turning off write-back caching will force SuperTrak66 to immediately save all writes to the drive.

- 1) To turn on/off write back caching from SuperCheck, double-click the Array icon in Tree View.
- 2) In the Information View, make sure Read Cache is turned "On".
- 3) Left-click the arrow next to the Write Cache selection box and choose "Write Back". This will enable "write" and "write back" caching functions from SuperTrak66 (see p. 98).

Viewing Cache Memory Statistics

The usage performance statistics of the SuperTrak66 memory provides the admin/user with information about the effectiveness and efficiency of the cache. Such statistics may be used as the basis of adjusting the amount of cache memory, size of read/write cache, or disabling cache.

- 1) To view Cache Statistics, double-click the Array icon in Tree View.
- 2) In the Information View, the statistics are located in the lower part of the grey area. A description of the statistics is provided on p. 98.

Viewing Array Object icons

In the Object View, icons for the hard drives that are connected to the current array selected will appear. Clicking on these icons will display hard drive information (see p. 87).



Figure 78: Array Object View

Information View

The Information View displays information pertinent to the operation of the chosen array, as seen below.

The screenshot shows a software interface for managing a RAID array. At the top, there are fields for 'Name' (RAID5_ARRAY), 'RAID level' (RAID5(PARITY-DISTRIBUTED)), 'Block Size (KB)' (64), and buttons for 'Commit' and 'Reset'. Below this, array details are listed: 'Array Size' (38848 MB (37.938 GB)) and 'Array Status' (functional). A section titled 'Cache Policy Settings' contains four dropdown menus: 'Read cache' (ON), 'Write Cache' (Write back), 'Read ahead mode' (User), and 'Read ahead range' (32). At the bottom, a table provides 'Cache Statistics' with the following data:

	Read cache hit	Cache Access Read	3024
Write cache hit	953	Cache Access Write	955
Partial read cache hit	2067	RAID Access Read	2069
Partial write cache hit	2	RAID Access Write	101
Dirty Usage	0	Reset Cache stats	

Figure 79: Array Information View

The fields displayed on the previous page are defined as follows:

Name

This field contains the name of the array.

RAID Level

The “RAID level” text box contains the RAID level information of the array. Pressing on the arrow along on the right edge of the text box (operational only when creating a new array) displays a list of choices (RAID levels 0, 1, 3, 5, or spanning).

Block Size (KB)

The “Block size” reflects the stripe block size information used by this array. This value cannot be modified here.

Cache Policy Settings

The “Cache policy settings” section contains several fields related to the caching mode of the array as follows:

Read cache

This field shows the current read cache setting. The options are “ON” and “Cache OFF.” “OFF” disables the entire caching system, including write caching.

Write cache

This is the write cache mode. Pressing on the arrow on the right edge of the text box displays a list of choices: “Off” (turns off write and write back caching); “write back” (enables the write and write back caching).

NOTE: *Flush parameters of write and write back caching are set in the IOPX Information View; see p. 84 for more information. Read Ahead mode and Range settings are not available in this version.*

Cache statistics

The cache statistics are displayed, but cannot be modified from here.

Read cache hit

This field shows the number of read cache requests made to this particular array since the field was last reset.

Write cache hit

This field shows the number of write cache requests made to this particular array since the field was last reset.

Partial read cache hit

This field shows the number of partial read cache requests made to this particular array since the field was last reset.

Partial write cache hit

This field shows the number of partial write cache hit requests made to this particular array since the field was last reset.

Dirty usage

This field shows the number of dirty blocks which the SuperTrak66 controller is holding in memory. These blocks will be flushed out to disk according to flush parameters set in the IOPx Information View (see page 84).

Cache access read

This field shows the total number of read requests (full and partial) made to a particular array since the field was last reset.

Cache access write

This field shows the total number of write requests (full and partial) made to a particular array since the field was last reset.

RAID access read

This field shows the total number of read requests made to a particular array which did not utilize caching and required a disk I/O to the RAID engine.

RAID access write

This field shows the total number of write requests made to a particular array which did not utilize caching and required a disk I/O to the RAID engine.

Commit/Reset cache stats

This button resets all fields in the Cache statistics box. If changes are made to any of the "Cache Policy" fields, press the "Commit" button to save the new settings, or "Reset" to make the values revert to their saved settings.

NOTE: The new settings will NOT be enabled until the next system boot.

NOTES

RAID Setup

This chapter discusses the various types of RAID arrays, their applications, including performance and reliability considerations, and general maintenance information. Given the variety of applications for RAID systems, the various RAID levels each have advantages and disadvantages.

Determining what RAID level to use for your system is an important step. To help guide you, we have divided this chapter into four parts:

- *Application Scenarios* describe different types of RAID system usage and match them to appropriate RAID configurations.
- *RAID Implementation* is for those systems which don't match well with our given application scenarios. The section covers general principles of each possible RAID level, and what types of system usage would benefit from those levels.
- *RAID Management & Operations* covers various fault scenarios, rebuilding/synchronization of drives, and drive partitioning/formatting which may be encountered during typical use.
- *Configuration Ideas* cover organizational models of array systems that use the previously-documented array types. It also points out additional advantages of various options.

This is certainly not intended as a complete, nor definitive reference on RAID technology applications, but merely serves as a first step to understanding the principles and proper usage of RAID. Please see *Appendix A: Technology Background* for more details on the functionality of RAID systems.

6.1 Application Scenarios

This section describes typical application scenarios for a RAID system. You may use this information as a guideline to tailor your RAID system configuration to meet your specific needs.

Keep in mind that when dealing with striped arrays, a larger stripe block size typically benefits arrays that tend to receive more random/smaller I/O requests. This is because such a setting tends to minimize seeks for a small request down to a single I/O operation. In contrast, a small stripe block size setting typically benefits arrays which tend to receive more sequential/larger I/O requests. This is because such a setting maximizes throughput by employing the data transfers of all drives simultaneously.

NOTE: *Though the following configurations recommend different RAID levels, RAID 3 is recommended in most single-threaded (non-server) desktop PC operations.*

High Performance Desktop PCs Working With Large Files

For video editing or desktop publishing PC stations where large files are stored on the hard drive, the need for performance outweighs that of data security. Accessing such large files requires sequential data transfers where performance can be improved by using striping (RAID Level 0). An ideal setup would have the large data files stored on the striped array, and the operating system boot drive using a separate disk controller (i.e. onboard motherboard controller).

When defining such an array, remember that a small stripe block size (such as 64KB) typically benefits arrays which tend to receive more sequential/large I/O requests (such as those used in video editing) since it maximizes throughput by employing the transfer of all drives simultaneously. The more drives assigned to this array, the better the throughput potential will be. Bear in mind, there are no fault tolerant data protection features provided in a RAID 0 array. It is exclusively geared towards gaining data transfer performance.

Corporate Desktop PC Requiring Data Protection

Under Windows NT, users have the option of enabling fault tolerance through software. SuperTrak66 permits a bootable, hardware-based fault tolerant solution to secure data from drive failure. RAID Level 3 offers the largest data capacity while maintaining fault tolerance using parity information. RAID 0+1 is also recommended since it offers higher data throughput and dual fault tolerance though it is more expensive in terms of drives.

Windows NT File Server

Windows NT tends to work heavily with large swap files for virtual memory. Under such a heavy load, the swap file I/Os undergo intensive activity. If the user's data files and the system files are located on the same drive, then the system performance is typically degraded because of the swap files.

One solution is to create two (2) mirroring (RAID Level 1) arrays using four (4) drives (2 drives per array). Array 1 can be used for the bootable system files. Array 2 can be used for the user or data files. With SuperTrak66, Array 1 would consist of drives 1 and 3, while Array 2 would use drives 2 and 4. Such a mirroring array architecture effectively boosts the performance of each array as well as providing data redundant data storage. In this configuration, heavy swap file access on Array 1 under significant load will not affect the transfer performance of the user/data array (Array 2).

Mid-Sized Windows NT Application Server

As previously described, it is best to separate the system/boot drive from the user/data drive because of the intensive file swapping I/Os which the disk system is subjected to from normal operation of Windows NT under heavy loads. For an NT application server, there are three possible recommended configurations for performance and reliability:

- 1) **Striping/Mirroring (RAID 0+1)** - 4 drives total. Two (2) drives are striped together and mirrored to two other striped drives. This provides good performance for read/write operations as well as single drive fault tolerance.
- 2) **Block and Parity Striping with "Hot" Spare (RAID 5)** - 4 drives total: 3 drives striped with rotating striped parity, and 1 drive unassigned as a "hot spare". This provides performance similar to configuration #1 above, but also offers better drive fault tolerance with immediate drive rebuilds.
- 3) **Block and Parity Striping (RAID 5)** - 4 drives total: all drives striped with rotating striped parity. This provides excellent read performance and single drive fault tolerance.

Configuration #1 exhibits good reliability from fault tolerance and good data transfer performance. Configuration #2 yields excellent reliability and good data transfers (though slower on writes than #1). Configuration #3 offers excellent performance and good reliability. The choice of which configuration to use should be made based on your particular needs to balance performance with data protection.

Largest Storage Capacity Required

In the event that the largest capacity possible is required, one of two possible configurations are suggested:

- 1) **Spanning** - This method effectively takes advantage of the total capacity of all attached drives. It essentially links all drives together to form one large drive regardless of the size of any individual drive (i.e. a 14.0GB drive and a 16.8GB drive will yield a 30.8GB storage array). This method offers no other RAID performance or data redundant features. It simply offers storage capacity.
- 2) **Striping (RAID 0)** - This method may be used with up to four (4) identical drives and *may* provide an additional performance gain over Spanning. If you connect four 16GB drives , you will get a 64GB drive array.

6.2 RAID Implementation

This section describes the advantages and disadvantages of the various RAID levels under particular workloads. This information can be used to select a RAID level appropriate for a system with known performance requirements which do not match any of the application scenarios in section 5.1.

Striping (RAID 0)

Striping is an effective method of gaining performance, but offers no fault tolerance. The measurable performance gain from using a striped array may vary depending on the typical application usage.

For applications which tend to perform random read/writes of varying sizes to the drive, a large stripe block size is recommended. With a large block size, a higher percentage of random read/write I/Os will be completed with just a single physical drive I/O. In such a scenario, four drives striped may provide up to four times the performance in multi-threaded applications versus a single drive.

For applications that tend to perform sequential/streaming read/writes (such as working with large database or project files, or editing audio/video), a small stripe block size is recommended. With a small block size, the extended sequential read/writes will assure that all physical drives are feeding data simultaneously. In such a scenario, four drives striped will provide much higher performance than a single drive. With the same small stripe block size, random multi-threaded read/writes of small-to-medium sized files may exhibit poorer performance when compared with a large block size.

Depending on the typical usage of the array, performance gains are directly affected by the stripe block size. Note that whatever block size you do choose, it will always be a compromise.

Mirroring (RAID 1)

Mirroring is a widely-used method of fault tolerance. Mirroring exactly duplicates the content of one drive on to the other drive for every write operation. If either drive fails, the array continues to function using the remaining working drive. This also allows for time to "hot" swap the failed drive and rebuild data on to a replacement drive without downtime (the RAID system can also invoke a "hot" spare drive, if available).

Generally, mirrored RAID arrays provide little-to-no performance enhancement. The SuperTrak66 system provides some performance enhancement on mirrored arrays by using elevator seek and load-balancing to achieve the fastest response possible from each access. Still, mirroring is not a performance replacement for striping. Use a mirrored array in environments which are not performance intensive, but rather must act as secure, reliable data/file servers.

Striping/Mirroring (RAID 0+1)

Striping/mirroring gives the best of both worlds for RAID 0 and 1. You've got all the performance gains and concerns of striping, coupled with all the advantages of data protection, elevator seek, and load balancing of mirroring. This type of array should be used for those cases which must match the application requirements of both striping and mirroring.

An added gain for RAID 0+1 is that, in certain situations, such an array offers double fault tolerance. Double fault tolerance allows your data array to continue to operate should two drives fail. We'll describe the four scenarios in which two drives can fail and still operate the array below. Under such an array structure, you would have a striped Array A consisting of two drives 1 and 2 (we'll call the drives, A1 and A2, for our purposes here). This striped pair of drives is duplicated by a second striped array B consisting of two drives 1 and 2 (we'll call them B1 and B2). The array will continue to operate if two drives fail as follows:

Scenario 1: Drives A1 and B2 fail, drives B1 and A2 operate the array.

Scenario 2: Drives A2 and B1 fail, drives B2 and A1 operate the array.

Scenario 3: Drives A1 and A2 fail, drives B1 and B2 operate the array.

Scenario 4: Drives B1 and B2 fail, drives A1 and A2 operate the array.

So long as there is a "1" drive from either "A" or "B" and a "2" drive from either "A" or "B", the array will remain functional. There are two double fault drive failure scenarios that will take the array offline. Should both drives A1 and B1 go down, or drives A2 and B2 go down simultaneously, the array will be offline.

Data striping with dedicated parity drive (RAID 3)

RAID Level 3 stripes data across several drives. Parity data is calculated and stored on a single dedicated parity drive. The parity information allows recovery if any single drive fails. Performance of a RAID 3 array is nearly the same as RAID 0 during data reads. However, during data writes, parity data must be calculated and updated to the dedicated drive each time. This slows performance during smaller

random writes (such as e-mail file servers, for example). For larger writes or sequential writes, the performance is still fairly fast. Because only one drive in the array stores redundant data, the cost per megabyte of a level 3 array is fairly low. RAID 3 requires hardware support for most efficient operation such as the processor employed by SuperTrak66 because of parity calculation.

Block and parity striping (RAID 5)

Block and parity striping (RAID 5) is useful for the performance gains offered from a striping array, with the added security feature of single fault tolerance. RAID 5 uses all drives to stripe data and also stripe parity data. Since parity is calculated on-the-fly during write operations, RAID 3 (which uses a dedicated drive for parity) experiences a performance bottleneck when the system constantly writes parity data to a single drive. Under RAID 5, parity is striped across all drives. Parity calculation and data writes are much faster, keeping all the drives in the array busy. This vastly improves random write performance.

In a nutshell, block and parity striping is best suited for those who wish to enjoy the performance advantages of a striping array, but who require at least single drive fault tolerance.

Spanning

In some instances, spanning may be a preferred model of combining drives together rather than striping. One reason that an array might be configured as spanning instead of striping would be to make full use of all the capacities of all drives in the array -- even if drives are mismatched in size. With striping, the array size is restricted to the number of drives times the storage capacity of the smallest drive. This cuts off portions of any drives that are larger than the smallest.

The other reason that spanning might be considered over striping may be performance. With striping, the performance is affected directly by the stripe block size. Block size should be tailored to the typical I/O on the drive -- whether it is generally more random or sequential. However, what if there is no predictability of the type of I/O access? What if both random and sequential I/Os occur unpredictably? The performance of a striped array will fluctuate. With spanning, the performance factor simply reflects a single drive's performance level, offers a more predictable transfer rate, and allows the use of mis-matched drives.

6.3 RAID Management & Operation

This segment covers issues pertaining to general RAID system management, operation, and terminology.

Critical & Offline Arrays

A fault tolerant array goes "critical" when a drive is removed or fails. Due to the fault tolerance of the array, the data is still available and online. However, once the array goes critical, the array has lost its fault tolerance, and performance may be adversely affected. If the fault was caused by a drive that was removed, the drive should be replaced by another drive, either identical or larger, in order for the RAID system to rebuild and restore optimal configuration.

A non-fault tolerant array goes "offline" when a drive is removed or fails. Since the array is not fault tolerant, the data stored in the array is no longer accessible. If the drive was removed, then it should be replaced to restore accessibility to the array. If the drive failed, then the entire array must be deleted and re-initialized since all data is considered lost.

Rebuilding/Synchronizing Fault Tolerant Arrays

Though a critical array can continue storage operations, it no longer offers fault tolerance and should be addressed as soon as possible by replacing the missing or failed drive(s).

Rebuild takes a replacement drive, assigns it to the array, then writes the redundancy data to it. Once the rebuild process is complete, the array status is upgraded from "critical" to "functional" and fault tolerance is restored.

Synchronization is a preventative maintenance measure used to avoid problems with data integrity. Synchronizing simply recalculates redundant data (similar to the rebuild process) and matches the data on the drive(s). Synchronization is a process that may be manually initiated or can be scheduled. You may choose to synchronize your array once per hour, day, week, or month.

Partitioning & Formatting Arrays

Like any other type of fixed disk media in your system, a RAID array must be partitioned and formatted before use. Depending on the O/S you use, there may or may not be various capacity limitations applicable for the different types of partitions. Use the same method of partitioning and formatting on an array as you would any other fixed disk.

6.4 Configuration Ideas

This part of the chapter covers additional ideas which are helpful to bear in mind while considering RAID configurations.

Fault Tolerance

Hard drives aren't the only things that can fail in a server room full of equipment, PCs, hard drives, cables, connections, and power supplies. There are literally hundreds of things that *could* go wrong on any given day. Although this is a disturbing thought that has led some to seriously question their career in computers and networking, there are some additional ways to prevent a server's RAID storage downtime and perhaps extend one's livelihood.

Since SuperTrak66 maintains four independent data channels and assigns just one drive per channel, users are protected in case a single drive OR an entire data channel fails. This provides an added level of protection than other RAID controllers that allow connecting two (more) drives to the same channel connector.

Hot Spare Drive(s)

A good precaution to protecting your array integrity in the event of drive failure is maintaining a hot spare drive. A hot spare is a drive which is connected to the array system, but is not assigned as a member of the array. In the event of the failure of a drive within a functioning fault tolerant array, the hot spare is activated as a member of the array. The spare drive effectively takes the place of the failed drive and the RAID system immediately begins to rebuild data on the drive. When the rebuild is complete, the array is returned to fault tolerant status. Once the failed drive is replaced, the new drive is automatically recognized as a hot spare in the event of a subsequent drive fault.

Some possibilities for making use of hot spares can include:

- A 2-drive mirrored array with 2 hot spare drives. This allows for triple drive fault tolerance before the array would be set offline.
- A 3drive RAID 3 array with 1 hot spare for double fault tolerance under RAID 3.
- A 3drive RAID 5 array with 1 hot spare for double fault tolerance under RAID 5.

Troubleshooting

This chapter provides information on how to interpret error conditions as reported by the SuperTrak66. This includes buzzer alarms, LED display codes, SuperTrak BIOS Error Messages, and Installation & Runtime Problems. The information contained here covers the most common error conditions that occur with SuperTrak66 installations.

7.1 Buzzer Alarms

Audible alarm indicators are triggered by various error events. The alarm itself may be constant or consist of a series of pulsed beeps.

Event: Buzzer occurs immediately from power-on, will not disengage, and the monitor will not turn on.

Cause: This condition may be caused by a bad copy of IxWorks in the flash RAM of the SuperTrak66 card. Return the SuperTrak66 card for a replacement. The IxWorks cannot be flashed by the user.

Cause: No memory exists or an unsupported memory type is installed on card. Check to ensure that the memory module is properly seated.

Event: Pulsing alarm occurs after the SuperTrak BIOS loads, accompanied by some error messages.

Cause: This is an indication of a problem with one or more drives in a RAID array under SuperTrak66 control. Look into the error message(s) which are displayed by the BIOS.

Event: Pulsing alarm begins when the SuperCheck utility accesses information on the drive housing.

Cause: This is an indication of a problem with a housing sensor. Either the operational temperature is out of nominal range, a cooling fan is not operational or there appears to be a power failure.

Event: Pulsing alarm begins when the SuperCheck utility is running.

Cause: This is likely an indication of an array problem. Check the status of all the arrays which are being monitored by the SuperCheck utility to verify their status.

7.2 LED Display Codes

LED displays appear on the backplane of the SuperTrak66 controller.

NOTE: *The LEDs on the backplane of the SuperTrak66 controller go through various patterns throughout the power-on and initialization process. However, only three different conditions remain constant.*

Event: LEDs on the card backplane turn on a few seconds after system power-on and show a pattern of ON-OFF-OFF-ON (1001).

Cause: This is the “normal” condition of the LEDs after power-on indicating that the IxWorks O/S loaded properly on the controller, and that the cache memory SIMM is functioning. This pattern requires no special attention.

Event: LEDs on the card backplane turn on a few seconds after power-on and move in a back and forth pattern.

Cause: This is the operational mode of the LEDs after the card has been initialized at power up. This indicates that the HDM thread is loaded and running correctly.

Event: LEDs on the card backplane turn on a few seconds after system power-on and show a pattern of ON-OFF-ON-OFF (1010).

Cause: This condition is caused by the SIMM memory on the SuperTrak66 controller. It is either not installed correctly, has problems with the connections, or may be an unsupported type of memory SIMM. Make sure the SIMM is seated correctly before trying an alternate memory module to correct the problem.

7.3 SuperTrak66 BIOS Error Messages

On bootup, the SuperTrak66 BIOS will initialize. Should an error be detected, the following messages will appear on screen and the bootup process halted.

Message: **No Array is defined...**

Cause: Either no drives are detected, or the drives detected are not assigned to an array. The SuperTrak66 BIOS only provides int 13h support for defined arrays. If there are no drives detected, or the detected drives are not assigned to an array, then the BIOS has nothing to support for int 13h. You must use the SuperBuild utility to define an array using drives attached to the SuperTrak66.

Message: **Problem is detected with Array : X**

Cause: There is a critical or offline array. A fault tolerant array will require one of the drives to be replaced (if a hot spare has not already been invoked). A non-fault tolerant array will be pulled completely offline, and is no longer accessible.

7.4 Installation & Runtime Problems

This segment covers certain problems which may be encountered during SuperTrak66 installation, configuration and run-time.

Event: SuperTrak66 BIOS does not show up

Cause: Being fully PnP compatible, the SuperTrak66 controller card offers no means of overriding the assigned memory address at which the BIOS will appear. It is entirely dependent upon the system BIOS being able to properly allocate a memory address resource for it. Make sure that your system PnP/PCI configuration allows for the BIOS to have a memory address between 640K and 1MB (more specifically, anywhere from 0xC8000 to 0xE8000, a total of 64KB). If your system BIOS supports ROM shadowing, you can disable ROM shadowing for a particular address at which the SuperTrak66 BIOS will then load.

Event: No LEDs on card light up

Cause: The SuperTrak66 card is not connecting with the computer. Make sure the card is fully seated in the PCI slot and that the data cable ground straps are properly secured.

Message: "NO ROM BASIC, SYSTEM HALTED"

Cause: This message is displayed when trying to boot a system from a drive which has no "active" partition on it. Use FDISK or a similar utility to set the active partition.

Event: Thermal recalibration interrupts sustained transfers

Cause: Thermal recalibration on older drives takes from 3 to 5 seconds depending on the drive. Newer drives have much better algorithms for thermal recalibration that last no longer than half a second. During this time, drive I/O may be buffered by the SuperTrak66 cache. Replacing older, slow recalibration drives with newer, faster models should be sufficient to sustain uninterrupted data transfers.

Technology Background

Introduction to RAID

RAID (Redundant Array of Independent Disks) allows multiple hard drives to be combined together to form one large logical drive or "array." As far as the operating system is concerned, the array represents a single storage device, and treats it as such. The RAID software and/or controller handles all of the individual drives on its own. The benefits of RAID can include: higher data transfer rates for increased server performance, increased overall storage capacity for a single drive designation (i.e. C, D, E, etc.), data redundancy/fault tolerance for ensuring continuous system operation in the event of a hard drive failure.

Different types of arrays use different organizational models and have varying benefits. The following outline breaks down the properties for each type of RAID array:

RAID 0 (Striping)

When a disk array is striped, the read and write blocks of data are interleaved between the sectors of multiple drives (see figure A1 next page). Performance is increased, since the workload is balanced between drives (or "members") that form the array. Identical drives are recommended for performance as well as data storage efficiency. The disk array's data capacity is equal to the number of drive members multiplied by the smallest array member's capacity.

For example, one 1GB and three 1.2GB drives will form a 4GB (4 x 1GB) disk array instead of 4.6 GB. The stripe block size value can be set logically from 1KB, 2KB, 4KB, 8KB, 16KB, 32KB, 64KB [Default], 128KB, 256KB, 512KB, or 1024KB. This selection will directly affect performance. Larger block sizes are better for random disk access (like email, POS, or web servers), while smaller sizes are better for sequential access. *NOTE: Promise recommends to use a block size smaller than 64KB for video-editing applications.*

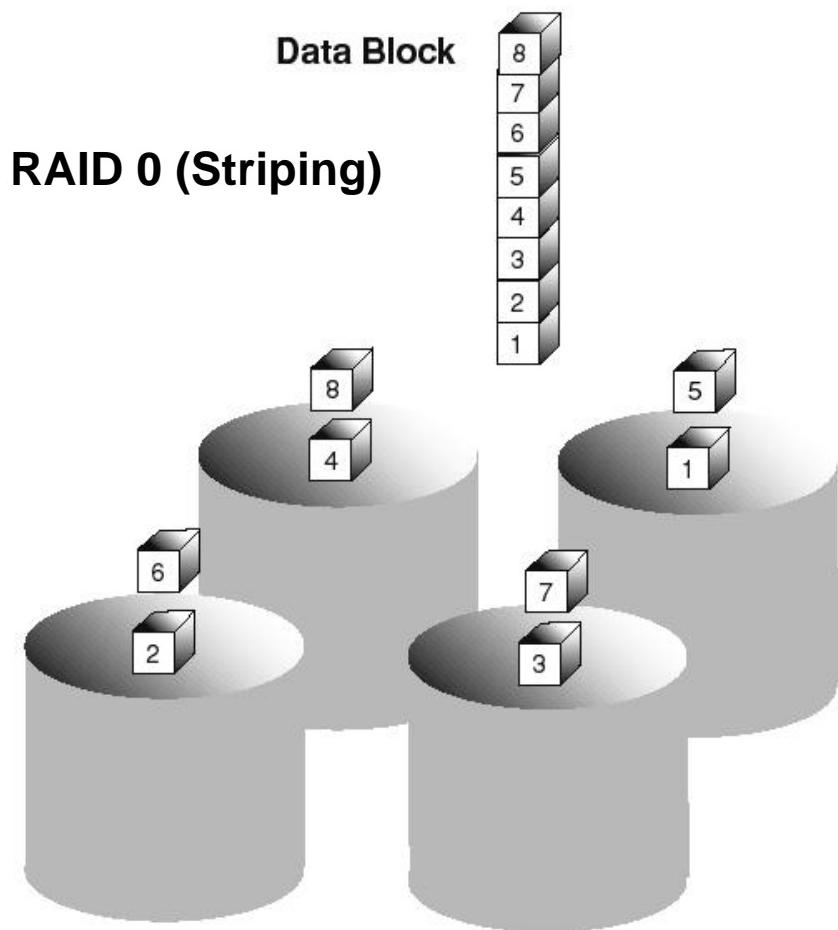


Figure A1: RAID 0 striping interleaves data across multiple drives

Mirroring (RAID 1)

When a disk array is mirrored, identical data is written to a pair of drives, while reads are performed in parallel. The reads are performed using elevator seek and load balancing techniques where the workload is distributed in the most efficient manner. Whichever drive is not busy and is positioned closer to the data will be accessed first.

Under RAID 1, if one physical drive suffers a mechanical failure or sector error, the other mirrored drive continues to function. This is called Fault Tolerance. Moreover, if a spare drive is present, the spare drive will be used as the replacement drive and data will begin to be mirrored to it from the remaining good drive.

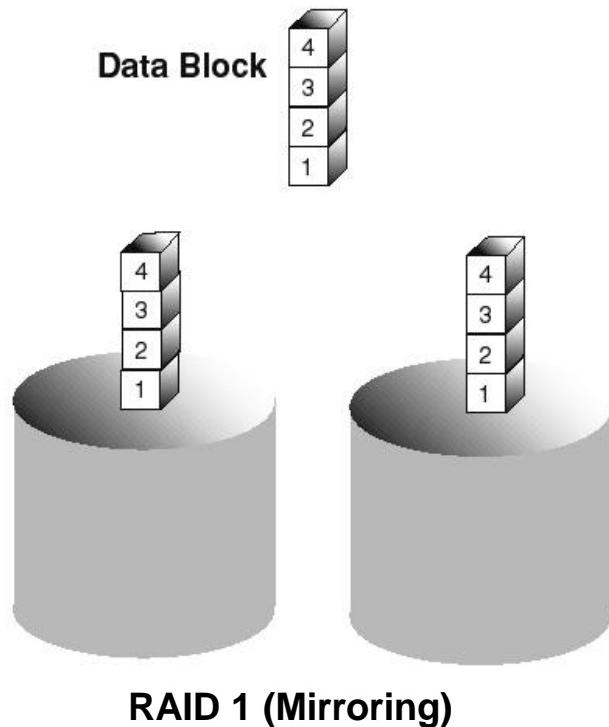


Figure A2: RAID 1 mirrors identical data to two drives

Due to the data redundancy of mirroring, the drive capacity of the array is only the size of the smallest drive. For example, two 1GB drives which have a combined capacity of 2GB instead would have 1GB of usable storage when set up in a mirrored array. Similar to RAID 0 striping, if drives of different capacities are used, there will also be unused capacity on the larger drive.

Striping/Mirror (RAID 0+1)

Striping/mirroring combines both of the previous array types. It can increase performance by reading and writing data in parallel while protecting data with duplication. A minimum of four drives are needed for striping/mirroring to be installed. With a four-drive disk array, two drives are striped together, and a second pair of striped drives are used to mirror the first pair of striped drives. The data capacity is similar to a standard mirroring array, with half of the total storage capacity dedicated for redundancy. An added plus for using RAID 0+1 is that, in many situations, such an array offers double fault tolerance. Double fault tolerance may allow your data array to continue to operate depending on which two drives fail.

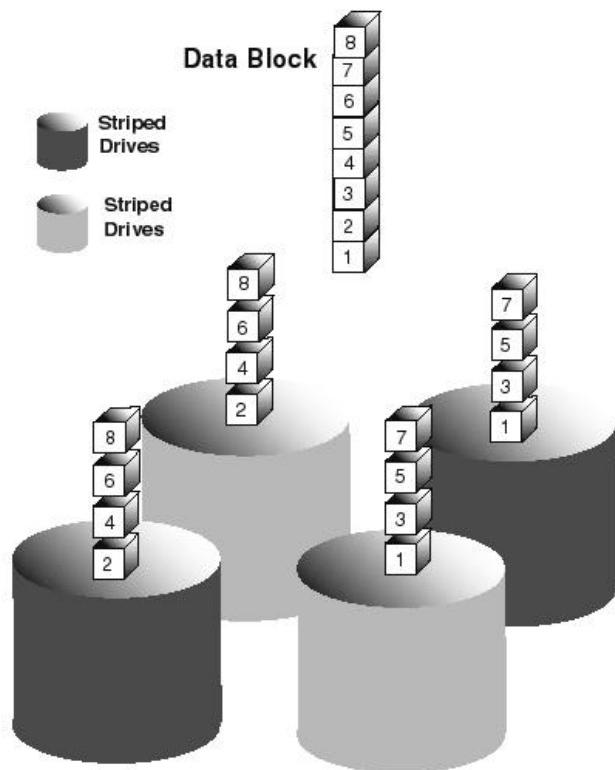


Figure A3: RAID 0+1 striping and mirroring of two drive pairs

Block Striping with Parity Drive (RAID 3)

RAID level 3 organizes data across the physical drives of the array, and stores parity information on to a drive dedicated to this purpose. This organization allows increased performance by accessing multiple drives simultaneously for each operation, as well as fault tolerance by providing parity data. In the event of a hard drive failure, data can be re-calculated by the RAID system based on the remaining drive data and the parity information.

The adjustable block size of the RAID 3 array allows for performance-tuning based on the typical I/O request sizes for your system. The block size must be set at the time it is created and cannot be adjusted dynamically. Generally, RAID Level 3 tends to exhibit lower random write performance due to the heavy workload going to the dedicated parity drive for parity recalculation for each I/O.

The capacity of a RAID 3 array is the smallest drive size multiplied by the number of drives less one. Hence, a RAID 3 array with (4) 1.0 GB hard drives will have a capacity of 3.0GB. An array with (2) 1.2GB hard drives and (1) 1.0GB hard drive will have a capacity of 2.0GB.

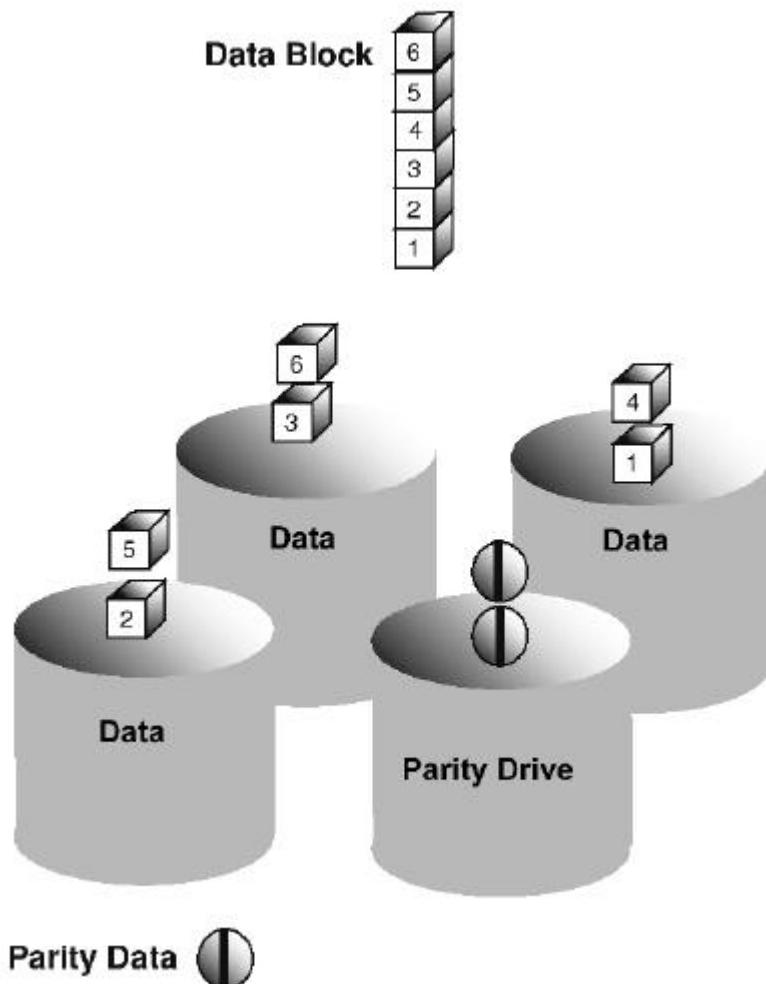


Figure A4: RAID 3 multiple drives stripe data w/ one dedicated parity drive

Block and Parity Striping (RAID 5)

RAID 5 is similar to RAID 3 as described above except that the parity data is rotated across the physical drives along with the block data. Having the parity data striped across all the physical drives in this manner removes the random write performance bottleneck of RAID 3. The total capacity of a RAID 5 array is calculated the same as a RAID 3 array.

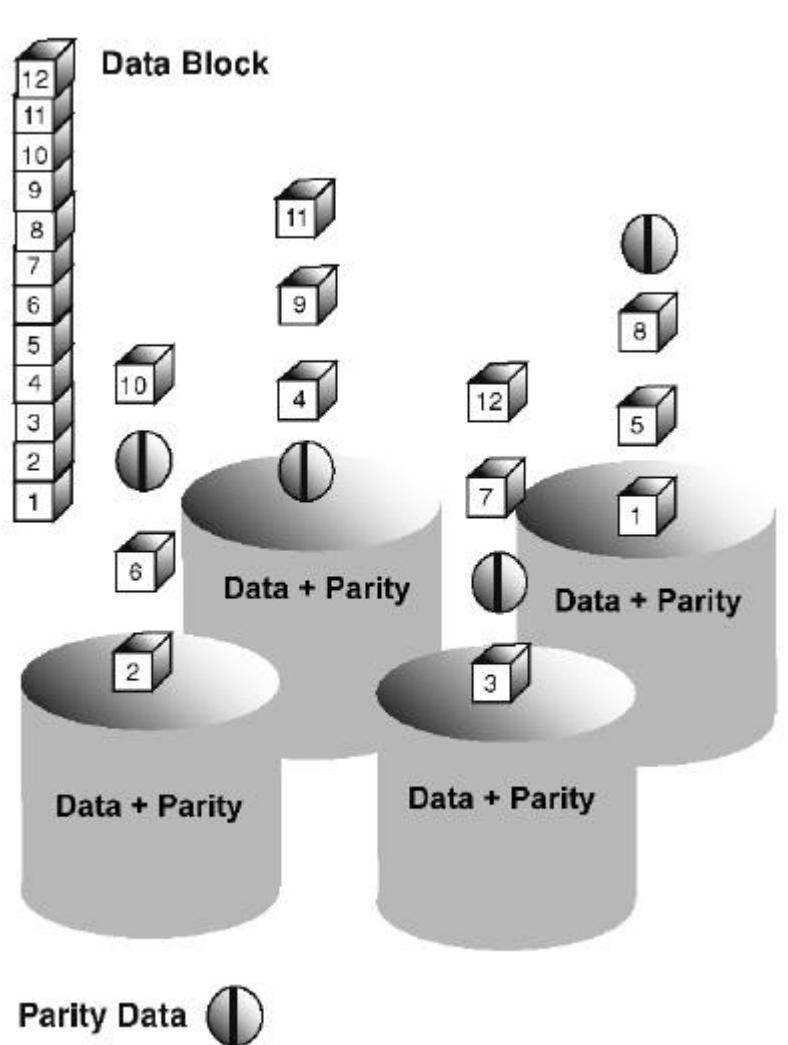


Figure A5: RAID 5 stripes all drives with data and parity info

Spanning

In a Spanning array, the disk array capacity is equal to the sum of all drives, even if the drives are of different capacities. Spanning stores data onto a drive until it is completely filled, then proceeds to store data on to the next drive in the array. There are no performance or fault tolerance array features in this type of array. When a disk member fails, the entire array is offline.

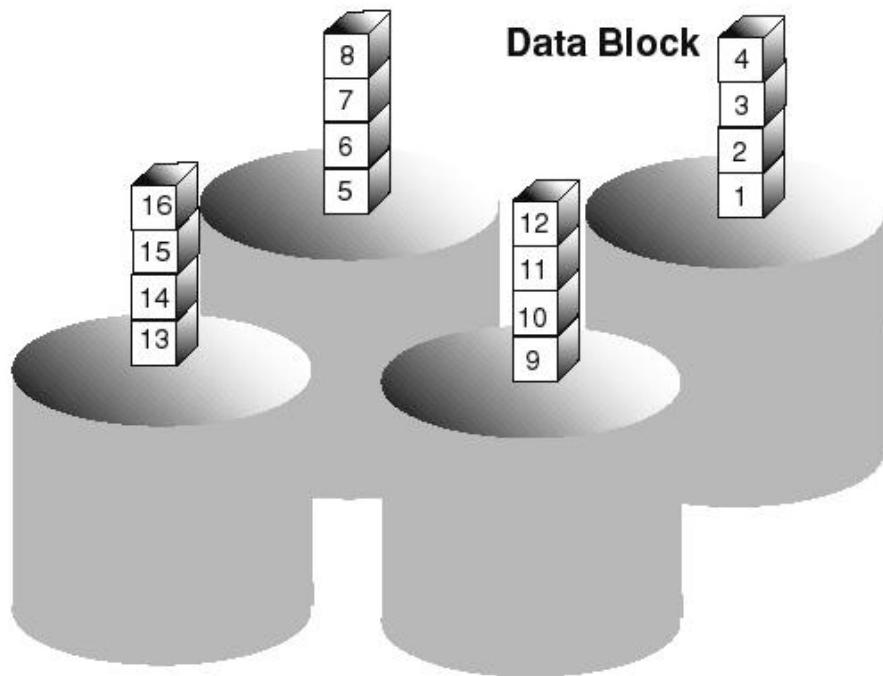


Figure A6: Spanning uses full capacity of drives

I₂O technology

I₂O, which stands for Intelligent Input/Output, addresses one of the weakest links in turning today's PCs into servers: degraded CPU performance due to heavy I/O processing. By offloading much of the I/O workload to the SuperTrak66 controller card processor, the I₂O architecture frees the system CPU from the many low-level requests involved in RAID operations.

I₂O is a forward-thinking approach to resolving the long-standing issue of platform interface independence. Since I₂O drivers are supplied by the maker of each operating system, I₂O product manufacturers can focus on product development rather than thinking about building specific drivers for each individual operating system. Theoretically, any product that is I₂O compatible becomes cross-platform-compatible with all operating systems that support I₂O.

At the time of this writing, only Windows NT 4.0 is officially supported by SuperTrak.

NOTES

SuperTrak66 Technical Specifications

Performance-Related Features

- True hardware-based RAID
- Single IRQ per RAID I/O
- Intelligent cache dynamically configures itself
- Support for up to 128MB (minimum of 8MB required) of EDO RAM in one 72 pin SIMM socket
- Microprocessor offloads tasks to reduce system CPU utilization
- User-configurable cache parameters for each array: dirty threshold, write policy, read ahead policy and flush policy
- User-configurable RAID stripe sizing allows controller to optimize operations for operating systems and applications
- User-configurable rebuild priority. Allows the user to keep RAID performance high while performing data rebuilds
- Multi-threaded command processing and disk operations
- Supports Tagged Command Queuing for optimal disk performance
- Advanced load balancing and elevator seek for optimal disk performance
- Hardware scatter gather DMA engine reduces CPU overhead
- Intelligent algorithm to optimize RAID 5 writes

Data Protection Features

- Support for RAID levels 0, 1, 0+1, 3, 5 and Spanning
- Meets all of the RAID Advisory Board's standards for RAID levels
- Automatic sector re-mapping allows defective hard disk sectors to be remapped with data recovery
- Automatic failed drive detection and transparent disk drive rebuild
- Onboard buzzer sound audible alarms in the event of failure
- Supports S.M.A.R.T. capable drives for predictive failure analysis messaging through SuperCheck GUI
- Hot swap capability
- Hot spare capability
- Configuration is stored on all disks in the array. This allows configuration to be restored to a new controller in case of controller failure

Host Interface

- On board 32-bit Intel i960RD RISC microprocessor
 - I₂O version 1.5 compliant
 - Supports up to 33MHz PCI bus speed
 - Supports up to 133MB per second burst data transfer rate across PCI bus
 - Provides 32-bit PCI bus mastering with 0 wait states.
 - PCI Local Bus version 2.1 compliant
 - Onboard 64K Flash BIOS
 - Firmware Flash EEPROM
 - NVRAM for error log for fast recovery
-

O/S support

- Windows NT 4.0+

Monitoring tools

- Monitoring utilities for Windows NT
- Remote monitor capable through Internet or LAN
- Notification of problems through email
- Monitors temperature, power and fan operation of enclosures, as well as status of drives and array
- Enclosure LEDs give drive and array status

Physical & environmental

- SuperTrak66 card size: 12" x 4.75" (30.5cm x 12.1cm)
- Operating temperature: 32 - 122°F (0 - 50°C)
- Relative humidity: 5 - 95% non-condensing
- Voltage: 5V±5%

Frequently Asked Questions

This section lists frequently asked questions involving pre-installation, drive issues, installation, and post-installation.

Pre-Installation (Speed, Device Types, Capacity, Cabling)

Q: What kind of hard drives can I use for a SuperTrak66 array?

A: You can use any Ultra ATA/66, Ultra ATA/33, or EIDE hard drive(s) to create arrays on the SuperTrak66. You should use matching drives for multiple-drive arrays to maximize capacity usage as well as performance. Ultra ATA/66 drives are recommended for highest performance.

Q: Can I run the SuperTrak66 controller at greater than 33Mhz PCI bus speed?

A: The current PCI 2.1 specification is for 33Mhz PCI bus speed. The SuperTrak66 controller is designed around this specification. In most cases, a higher PCI bus speed will result in a variety of different errors. Some users have been able to get higher speeds to work. However, since such usage is out of specification, we do not support it. Any damage resulting from such implementation will void the warranty.

Q: Can I use ATAPI devices on the SuperTrak66 controller?

A: No. There is no driver layer on the SuperTrak66 controller which will support ATAPI packet messages.

Q: Can I substitute any other type of Ultra ATA/66 or IDE cables for SuperTrak66 cables if I need to use a longer data cable?

A: No! The SuperTrak66 data cables are custom-tailored to Ultra ATA/66 data transfer requirements with special shielding and a maximum IDE specification length of 18 inches.

Q: Does the SuperTrak66 controller work with my system's APM?

A: No. APM is based on IDE or SCSI drives. Since SuperTrak66 uses the IO platform-independent interface, the system does not see drives attached to the SuperTrak66 as either IDE or SCSI drives. Thus, it will not issue APM commands. Promise may or may not opt to implement some form of flexible power management for the SuperTrak66 controller in a future release.

Drive Issues

Q: Can I add a drive to a SuperTrak66 RAID array via hot-swap and dynamically adjust the array size/configuration?

A: No. The SuperTrak66 system does not support dynamically adjustable RAID size/configurations.

Q: Can I take a set of drives which make up an array created on one SuperTrak66 server and move it to another SuperTrak66 server?

A: Yes. All SuperTrak66 servers read the arrays the same way. Once the drives are all connected, you must restart the system for it to recognize the newly-inserted array configuration.

Q: Can I take a drive used in a SuperTrak66 array and access it directly with a different controller, such as the one integrated on the motherboard?

A: Yes, but only under certain configurations. The following array configurations will allow the drive(s) to be accessed individually on another controller: mirror (RAID 1), single drive striped (RAID 0), or single drive spanned array. Multiple drives striped or spanned will not work. Also, the other controller must address the drives as LBA, not CHS.

Q: If I have a problem with one of the drives on the SuperTrak66, how can I low level format it to correct the problem?

A: Do NOT do this. Low-level formatting IDE drives is unnecessary and generally does not correct problems which might be experienced during use. Errors such as bad sectors or ECC/CRC failure are best remedied by completely replacing the drive. For this reason, drives attached to the SuperTrak66 controller should NOT be low-level formatted.

Q: Do I have to install disk management software on my array in order to access the full storage capacity of drives?

A: No! The purpose of disk management software is to provide support for hard drives larger than 504MB to those systems which have no intrinsic support. Since the SuperTrak66 provides all the large disk support you need to access your array(s), this software is unnecessary, and in some cases may even cause accessibility problems.

Q: What system BIOS CMOS settings do I use for the drives on the SuperTrak66 controller?

A: None. The drives on the SuperTrak66 controller are supported by the SuperTrak BIOS and/or O/S drivers, not by your system BIOS.

Installation Issues (NT, I₂O Drivers, Drive Capacity, Booting, IRQ Settings)

Q: Why doesn't Windows NT see my SuperTrak66 controller?

A: You need to install the Windows NT I₂O drivers to be able to see the card and the attached drives (arrays).

Q: Why are some drives recognized by the SuperTrak66 Array Setup utilities with only partial capacity?

A: Some hard drive models are shipped with a jumper that reduces the addressable capacity of the drive. This is to prevent problems with older systems which won't support larger drives. Consult the documentation accompanying the hard drive to set the jumper appropriately in order to utilize the full capacity of the drive.

Q: Can I make an array on the SuperTrak66 bootable?

A: Yes. The SuperTrak66 BIOS provides minimal system-level support for the array as an int 13h drive. Use the BIOS SuperBuild setup utility to create an array in the Array 1 slot position. Then partition and format the new array accordingly for the O/S to boot. Make sure the boot sequence assigns the SuperTrak66 array as the boot device if you have drives running from the motherboard or another controller card.

Q: How can I change the system boot sequence in order to boot from the SuperTrak66 array or SCSI?

A: The boot sequence is controlled by the system BIOS. As far as the system BIOS is concerned, the SuperTrak66 controller and attached arrays are categorized as a "SCSI" device (even though the BIOS will not attempt to access the SuperTrak as any type of SCSI device). This allows you to set the boot sequence in your BIOS CMOS setup utility to boot from "SCSI" (an add-in controller like the SuperTrak66) first, rather than "IDE" (an IDE controller built onto the motherboard or one which effectively replaces it). If there are multiple SCSI add-in controllers in the system, then the boot sequence among them will be determined exclusively by their PCI slot priority. PCI slot #1 will be first, slot #2 second, etc. Put the SuperTrak66 controller in the PCI slot where it will be accessed ahead of other SCSI controllers if you want to boot from the array.

Q: How can I change the IRQ setting for the SuperTrak66 controller?

A: The SuperTrak66 controller is fully PCI PnP. This means all the resources that it uses are assigned by the PCI BIOS on the motherboard. While SuperTrak66 supports the IRQ sharing feature, this will only work if all the PCI devices used support it. If your motherboard allows you to control the assignment of these resources, you may be able to change them manually (see motherboard BIOS documentation for how to perform this task). If not, resource assignments may otherwise be changed by switching the PCI slot that the controller is in.

Post-Installation**(Power Up, BIOS Initialization, Memory Managers, Caching, Booting)****Q: Why does my SuperTrak66 take so long to power up and get a signal to the VGA monitor now that the SuperTrak66 card is installed?**

A: The process should only take a few seconds at most. The SuperTrak66 controller performs an intensive memory check each time at power-up. This assures the integrity of your system data during normal data caching operation. The greater the amount of RAM installed on the controller, the longer this process takes to verify memory. The memory scan places a hold on the PCI bus and suspends system boot temporarily until the scan is complete.

Q: Why does the SuperTrak66 BIOS take so long to initialize the card? Is there a problem?

A: Generally, this does not indicate a problem. The SuperTrak66 card initialization is a rather involved process. The BIOS sends a few messages to the SuperTrak66 controller during the initialization and spends most of the time waiting for the card to complete the requested task. The tasks initialize the hardware device layer (which performs hard drive detection), the RAID engine (which detects and sets up multiple tables and buffers for RAID operation), the cache module, device attachment messages, and system control information (including controller card, driver and O/S version control information). The amount of time required for initialization is not an indication of a problem. If there are any problems, you will be alerted immediately.

Q: How can I be sure that write-back cache has flushed before I reboot after partitioning and formatting an array?

A: A cache flush is always triggered immediately following any write from the int 13h BIOS. This level of array support is what FDISK and FORMAT use to access the drive from MSDOS. With no delay in the write operation, there is no cached data waiting to be written. If you are running from NT or other O/S, as usual, make sure to shut down NT properly so that the O/S can trigger and flush the cache as necessary. *NOTE: powering off before a proper shutdown is complete may result in the loss of unwritten data in cache, possibly causing problems in a subsequent boot or access to the data which was unwritten.*

Q: Why can't I see the drives on the SuperTrak66 under FDISK?

A: If you have not created an array, the physical drive(s) attached to the SuperTrak66 card will not be recognized by the operating system. The SuperTrak66 controller is dedicated to RAID array management and does not provide any means of addressing individual hard drives through the int 13h interface used by FDISK. In order to access drives on the SuperTrak66 from MSDOS at all, you must first create a RAID array.

Q: Why can't I see the array I just created with SuperCheck under the NT Disk Administrator?

A: Since Windows NT does not yet support any method of dynamically adding and/or removing logical devices to/from the system, you must restart NT. The next boot will show the new array under Disk Administrator.

NOTES

Contacting Technical Support

Promise Technical Support provides several support options for Promise users to access information and updates. We encourage you to use one of our electronic services, which provide product information updates for the most efficient service and support.

If you decide to contact us, please have the following information available:

- Product model and serial number
- BIOS and driver version numbers
- A description of the problem / situation
- System configuration information, including: motherboard and CPU type, hard drive model(s), IDE/ATAPI drives & devices, and other controllers.

Technical Support Services

Promise Online TM	http://www.promise.com
Web Site	(tech documents, drivers, utilities, etc.)

USA Tech Support Center

E-mail Support	support@promise.com
Fax Technical Support	(408) 452-9163 Attention: Technical Support
Phone Technical Support	(408) 452-1180 8:30-5:00pm M-F Pacific Standard Time
If you wish to write us for support:	Promise Technology, Inc. Attn: Technical Support 1460 Koll Circle, Suite A San Jose, CA 95112 USA

European Tech Support

E-mail Support	support@promise.nl
Fax Technical Support	+31 (0) 40 256 94 63 Attention: Technical Support
Phone Technical Support	+31 (0) 40 256 94 61 8:30-5:00pm The Netherlands Time
If you wish to write us for support:	Promise Technology Europe B.V. Attn: Technical Support 1European Buisness Centre, Unit 1.25 Luchthavenweg 81 5657 EA Eindhoven, The Netherlands

Pacific Rim Sales Office

E-mail Support	support@promise.com.tw
Fax Technical Support	+886 3 578 23 90 Attention: Technical Support
Phone Technical Support	+886 3 578 23 95 9:00-6:00pm Taiwan Time
If you wish to write us for support:	Promise Technology, Inc. Attn: Technical Support 4F, No.1, Prosperity 1 st Road Science-based Industrial Park Hsinchu, Taiwan

China Office

E-mail Support	support-china@promise.com
Fax Technical Support	+86 (0) 10 6872 3940 Attention: Technical Support
Phone Technical Support	+86 (0) 10 6872 3942 9:00-6:00pm China Time
If you wish to write us for support:	Promise Technology China Attn: Technical Support Room 3217, No. 15, Bai Shi Qiao Road Hai Dian District Beijing 100081 P.R. China

Limited Warranty

Promise Technology, Inc. ("Promise") warrants that for one (1) year from the time of the delivery of the product to the original end user:

- a) the product will conform to Promise's specifications;
- b) the product will be free from defects in material and workmanship under normal use and service.

This warranty:

- a) applies only to products which are new and in cartons on the date of purchase;
- b) is not transferable;
- c) is valid only when accompanied by a copy of the original purchase invoice.

This warranty shall not apply to defects resulting from:

- a) improper or inadequate maintenance, or unauthorized modification(s), performed by the end user;
- b) operation outside the environmental specifications for the product;
- c) accident, misuse, negligence, misapplication, abuse, natural or personal disaster, or maintenance by anyone other than a Promise or a Promise-authorized service center.

Disclaimer of other warranties

This warranty covers only parts and labor, and excludes coverage on software items as expressly set above.

Except as expressly set forth above, Promise disclaims any warranties, expressed or implied, by statute or otherwise, regarding the product, including, without limitation, any warranties for fitness for any purpose, quality, merchantability, non-infringement, or otherwise. Promise makes no warranty or representation concerning the suitability of any product for use with any other item. You assume full responsibility for selecting products and for ensuring that the products selected are compatible and appropriate for use with other goods with which they will be used.

Promise does not warrant that any product is free from errors or that it will interface without problems with your computer system. It is your responsibility to back up or otherwise save important data before installing any product and continue to back up your important data regularly.

Promise's sole responsibility with respect to any product is to do one of the following:

- a) replace the product with a conforming unit of the same or superior product;
- b) repair the product;
- c) recover the product and refund the purchase price for the product.

Promise shall not be liable for the cost of procuring substitute goods, services, lost profits, unrealized savings, equipment damage, costs of recovering, reprogramming, or reproducing of programs or data stored in or used with the products, or for any other general, special, consequential, indirect, incidental, or punitive damages, whether in contract, tort, or otherwise, notwithstanding the failure of the essential purpose of the foregoing remedy and regardless of whether Promise has been advised of the possibility of such damages. Promise is not an insurer. If you desire insurance against such damage, you must obtain insurance from another party.

Some states do not allow the exclusion or limitation of incidental or consequential damages for consumer products, so the above limitation may not apply to you.

This warranty gives specific legal rights, and you may also have other rights which vary from state to state. This limited warranty is governed by the State of California.

Returning Product For Repair

If you suspect a product is not working properly, or if you have any questions about your product, contact our Technical Support Staff through one of our Technical Services, making sure to provide the following information:

- Product model and serial number (required);
- Return shipping address;
- Daytime phone number;
- Description of the problem;
- Copy of the original purchase invoice.

The technician will assist you in determining whether the product requires repair. If the product needs repair, the Technical Support Department will issue an RMA (Return Merchandise Authorization) number.

Return only the specific product covered by the warranty (do not ship cables, manuals, diskettes, etc.), with a copy of your proof of purchase to:

Promise Technology, Inc.
Customer Service Dept.
Attn.: RMA # _____
1460 Koll Circle
San Jose, CA 95112

You must follow the packaging guidelines for returning products:

- Use the original shipping carton and packaging
- Include a summary of the product's problem(s)
- Write an attention line on the box with the RMA number
- Include a copy of proof of purchase

You are responsible for the cost of insurance and shipment of the product to Promise. Note that damage incurred due to improper transport or packaging is not covered under the Limited Warranty.

When repairing returned product(s), Promise may replace defective parts with new or reconditioned parts, or replace the entire unit with a new or reconditioned unit. In the event of a replacement, the replacement unit will be under warranty for the remainder of the original warranty term from purchase date, or 30 days, whichever is longer.

Promise will pay for standard return shipping charges only. You will be required to pay for any additional shipping options (such as express shipping).

Your Responsibilities

You are responsible for determining whether the product is appropriate for your use and will interface with your equipment without malfunction or damage. You are also responsible for backing up your data before installing any product and for regularly backing up your data after installing the product. Promise is not liable for any damage to equipment or data loss resulting from the use of any product.

Glossary of Terms

API	Application Programming Interface: A software layer which provides a common method of accessing several types of hardware or other software that it supports. APIs reduce the technical expertise needed for applications to be able to support a wide array of hardware devices and software packages directly.
array	As used in RAID, an array consists of one or more hard drives which are logically combined to form a single storage drive. Arrays are categorized by the method in which they are accessed to logically organize data on them for purposes of performance enhancement, capacity augmentation, and/or data redundancy protection.
ATA	AT Attachment: A standard for interfacing peripheral devices to "PC/AT compatible" machines.
BIOS	Basic Input/Output Set: Typically a set of system instructions stored in ROM or NVRAM that contain a rudimentary foundation for system initialization and operation. These programs provide just enough support to allow an O/S to load and take over the PC operations, or to provide a backward-compatible API to interface with new types of devices to a system otherwise unsuited to handle such devices.
bit	The fundamental representation of a numeric digit in the binary (base 2) numbering system. A bit consists of the value 1 or 0 (also referenced as "ON/OFF" and "TRUE/FALSE", respectively). By grouping bits together, the combinations of values that the bit groups represent increase exponentially, forming the basis of binary numbering. The most basic access to a collection of bits is a byte.
boot	A term describing the process of a computer system initializing itself and loading an O/S from a disk drive, network connection, or other source.
bootstrap	The first sector of a partition containing executable machine code information. The machine code contained here provides enough instruction to tell the system how to begin loading the rest of the O/S from disk to memory. If there is no O/S installed on a drive, this sector usually contains code which prints a message "Insert system disk.." or similar. Often, a program called a "boot manager" may be installed onto a drive which installs itself to give the operator further O/S loading options from various drive partitions in the system.
byte	The most basic data chunk that is directly accessible to i80x86 microprocessors (and many others). A byte consists of exactly 8 bits.

cache	A method of enhancing disk system performance by reducing and advancing physical drive accesses in an intelligent manner, by temporarily storing disk data in local RAM. Subsequent access to cached data is then fetched from RAM which is much faster than accessing the physical media. Using cache memory, systems can provide read ahead caching where more data is brought into cache memory than was requested by the O/S during a particular access. Read ahead caching anticipates the next read request by the O/S to increase data read response when the request "hits" the advanced cache. The same principle is applied with write back cache, the practice of holding disk write requests in cache memory until a flush is performed. This effectively combines multiple disk I/Os into a single operation which may otherwise have been fragmented by the O/S.
channel	Refers to a specific channel connection on the SuperTrak66 controller card. Each cable connected to the controller represents a single channel.
CHS	Cylinder Head Sector: An older method of defining the geometry and accessing the hardware of media devices such as hard drives. A drive which has multiple platters is divided up into cylinders which are in turn divided by heads. Each cylinder/head pair has a sector address up to the SPT value. Today, this method is used more for backward-compatibility since LBA is now the prevailing addressing method.
DDM	Device Driver Module: A software component of IxWorks. DDM includes both HDM and ISM, and (potentially) other types of modules.
dirty	A term used loosely to describe blocks of drive data in cache memory which require a flush.
DMA	Direct Memory Access: DMA is a system bus device which has its own logic for addressing a piece of memory in the system without the use of a CPU to move data. This allows I/O devices such as hard drive controllers to push data to/from memory without burdening the system CPU to perform the data transfer to the drive through a mechanism such as PIO.
EDO	Extended Data Out RAM: A type of RAM which has a larger internal working buffer, boosting performance under normal use.

EIDE	Enhanced Integrated Device Electronics: An I/O data bus model used by ATA and ATAPI devices. This model is superior to IDE allowing the addition of larger capacity devices through LBA addressing, and better performance through faster PIO and DMA data transfer specifications.
flash	A term used loosely to describe the ability to upload and download memory images to/from NVRAM in order to update some type of system software.
flush	A term describing the action of taking data written to cache memory and writing the data out to the disk.
GB	Gigabyte: 1,073,741,824 bytes. 1,024•1,024•1,024; commonly mistaken for 1,000,000,000 which is incorrect. 1,073,741,824 is derived from the base 2 (binary) numbering system.
geometry	A set of values which numerically describe the address limits of a media device, and enables software and hardware to fully control the device. CHS and LBA are types of geometry.
HDM	Hardware Driver Module: A software component of IxWorks which is used to access hardware devices attached through a local bus.
hot spare	A drive which remains attached to the RAID controller system, but is unassigned to any particular array. In the event of a drive fault in a fault tolerant array, the drive is activated to take the place of the failed drive.
hot swap	The ability to install/remove drive(s) from a system while it is still powered up and running. A system supporting hot swap should not have any problem detecting the removal or reattachment of a hot swapped device. SuperTrak66 takes advantage of hot swapping to allow removal and replacement of dysfunctional hard drives using the optional FastSwap66 or SuperSwap66 removable drive housings.
IDE	Integrated Device Electronics: An I/O data bus model used by ATA and ATAPI devices.
I/O	Input / Output: A term loosely used to describe the movement of data to and from hardware devices.
IP	Internet Protocol: Often referred to as "IP Address." IP is a series of 4 numbers separated by "." (ex. "127.0.0.1".) The IP represents a unique system on a network that may be communicated with. Each number may hold a value with ranges from 0 (zero) to 255. This translates to over 4 billion possible combinations of addresses.
iO	Intelligent Input/Output: A message-based I/O architecture.
ISM	Intermediate Software Module: A software component of IxWorks used to access devices hosted by a HDM as used by SuperTrak. The ISM is more intelligent and provides features like RAID and caching.

IxWorks	An O/S built for dedicated processors. In the case of SuperTrak, IxWorks uses an I ₂ O message frame architecture and i960 microprocessor.
JBOD	Just a Bunch Of Drives: Another term describing the spanning type of RAID arrays.
KB	Kilobyte: 1,024 bytes. commonly mistaken for 1,000 which is incorrect. 1,024 is derived from the base 2 (binary) numbering system.
LBA	Logical Block Addressing: LBA is an alternative form of sector addressing for accessing a media device. It is generally preferred over the older CHS method. With LBA, each sector on the device is assigned an address number (in order starting from 0) which increments for each sector. LBA addresses are typically stored as 32-bit values, occasionally they are 64-bit.
MB	Megabyte: 1,048,576 bytes. 1,024•1,024; commonly mistaken for 1,000,000 which is incorrect. 1,048,576 is derived from the base 2 (binary) numbering system.
MBR	Master Boot Record: The MBR is contained in the first physical sector of a hard drive compatible with modern PCs. It contains information about the partitions on the drive, and where the bootstrap information is located.
NVRAM	NonVolatile RAM: A type of memory that retains its contents after power is shut down and, later, when power is restored. Unlike ROM memory, data contained in NVRAM may be written to. The memory write is a complex (typically proprietary operation) which requires special software support. This prevents most applications from altering the contents of the NVRAM. This property makes NVRAM ideal for usage in flash memory implementations.
O/S	Operating System: A term commonly used to reference the CPU operating system on a platform station/server. Windows NT and Unix are examples of common operating systems.
pane	A segment of an application window, often encapsulated by some type of border, which contains a separate section of the application.

parity A mechanism for making the sum of two binary digits (bits) equal to a pre-specified digit. Such a system may use "even parity," in which case the sum is made even, or "odd parity," in which the sum is made odd. As an example if we had two pieces of data, 'A = 4' and 'B = 6', parity data would be defined as $A+B = 10$. In RAID 3 or 5 applications, such parity data can be used to reconstruct corrupt or missing data from a failed drive. From the example above, if hard drive 'A' has failed, and the only data we have is that on hard drive 'B' and parity data contained on another drive, the RAID controller will work backwards to rebuild the missing data on to a replacement drive. The data on hard drive 'A' can be calculated since ' $A = \text{parity} - B$ ' (or ' $A = 10 - 6$ '), yielding ' $A = 4$ '.

partial read cache hit

A partial read cache request has data which is partially in the cache and part of which must be read directly from the disk.

partial write cache hit

A partial write hit occurs when part of the data transfer is already in write back cache and part has to have new memory blocks allocated to it. Partial write hit requests usually result in additional disk operations because of the block allocation. By contrast, a full write hit would not require an additional disk operation for block allocation.

partition A logical division created on drive media, as recognized by the O/S. A drive with a single partition defines up to 100% of the drive as belonging to it, depending on the file system used to set up the drive. Multiple partitions may be used to break large drives into smaller chunks which may be more manageable. Each partition may be addressed by the O/S as a separate drive letter.

PIO Programmed Input/Output: A method of performing read/write operations on an IDE drive in a programmed timing interval between block requests. This type of I/O requires some sort of CPU utilization in order to handle the I/O and store the data transferred to/from the drive.

RAID Redundant Array of Independent Disks: A model of logically organizing multiple hard drives accessible by an O/S as a single logical device. Such an architecture of drives may improve I/O performance by accessing multiple drives simultaneously, and/or add fault tolerance by providing redundancy and parity configuration options.

RAM Random Access Memory: A type of memory which is used in all types of computer systems and by peripheral devices, accessible for read or write at any time. RAM loses its contents after a power-down.

ROM	Read Only Memory: A type of memory which is commonly adapted to computer systems in order to hold crucial data or programming code. Such information is critical to booting and/or operating the system's basic functions. ROM is read-only which means that the system may read information in from the memory and begin to use it, but cannot modify or erase the information. ROMs retain their data content without having power applied to them, and the same data is available every time the device is powered on. A system or peripheral BIOS are generally located in a ROM, or NVRAM.
SMTP	Simple Message Transfer Protocol: This is a network communications protocol used by several layers of network communications, including email.
SPT	Sectors Per Track: This field of the drive's CHS geometry defines the highest value for the sector.
striping	Refers to the storage of sequential block data across multiple drives in a RAID array group.
TCP/IP	Transfer Control Protocol / Internet Protocol: This is a software interface for drivers or APIs which provide networking services for a machine.
UDMA	Ultra DMA: A hard drive transfer specification with higher data rates than its predecessor, DMA.
XOR	eXclusive OR: This is a common, bit-wise logical operation performed with two inputs, yielding one output. In RAID applications, the known results of this operation allows for reconstruction of data supplied as input to XOR, and the output is used to generate parity.

INDEX

A

- Administration
 - Assigning user rights 75
 - Password protection..... 76
- Arrays
 - Administration of 41
 - Bootable** 15, 28
 - Creating .16, 22, 24, 26, 27, 44, 79
 - Critical & Offline 107
 - Deleting 16, 29, 45, 89
 - Deleting drives 89
 - Environments 102
 - Fault tolerance in RAID 1.... 45, 89
 - Hot Spare Drives 28
 - Partitioning/Formatting..... 3, 107
 - RAID 0..... 24
 - RAID 1..... 24
 - RAID 3..... 26
 - RAID 5..... 26
 - Rebuild Wizard 91
 - Rebuilding 90, 107
 - Recognition order 28
 - Saved information..... 29
 - Spanning..... 27
 - Synchronization** 73, 93
 - Synchronizing..... 107
 - Viewing status of,..... 89

B

- Beeper
 - Enable/Disable..... 87
- BIOS 15
 - Creating arrays 16, 22
 - Deleting arrays..... 16, 29
 - Driver version 30
 - Error messages 111
 - Start-up sequence..... 17
 - Using SuperBuild 17
- Block Size..... 102
 - Changing..... 24, 26, 27
- Creating arrays 16
- Performance issues..102, 104, 106
- Viewing 98

Bootable Array

- Assigning 28
- Buzzer Alarms 109

C

Cache

- Array info..... 98
- Controller info..... 84
- Flushing..... 82
- Installation 8
- Read 95
- Timing 81
- User rights..... 77
- View statistics 96, 98
- Write-back..... 95

CMOS Setup

- Boot sequence..... 112
- Drive parameters** 8
- Controller 2, 6
 - Attached components..... 79
 - Cache policy..... 83
 - Firmware info 83
- LED indicators** 110

D

Drivers

- Existing Win NT 4.x 12
- Existing Win NT3.5 12
- Removing Win NT 3.5x 12

E

E-mail

- Events list 46, 68
- Notification 41, 45, 67
- Receiver list 45, 67
- Removing user 46, 68

- Enclosure Management 87
 - Beeper functions 87

Errors

- Halt On 16
- Event Log
 - Clearing..... 80
 - Event Viewer..... 81
 - Viewing..... 80

H

Hard drives.....	See Arrays
Assignment of,..	21, 24, 25, 26, 27, 86
Hot Spare.....	103, 105, 108
Hot Spare.....	28
Hot Swap	105
Installation	8
Model info.....	86
View Assignment of,	21
Viewing	47, 85, 86
Hardware	
Cables	6
I2O support.....	78
Number of channels.....	30
Processor memory	30
Quick install	7
Step-by-step install	8
SuperTrak controller.....	6
Help Menu	
About.....	53
Finding topics	55
Viewing Autodemo	53
Viewing contents.....	54

I

I2O	
Version info	78
Information View	50

L

LED Indicators	
Controller backplane	110

M

Message Agent	2, 32, 44
Accessing	65
Message Server	2, 32, 43
Connecting	61
Creating	58
Deleting.....	60
Disconnecting.....	60
Switching Servers.....	62
Mirroring Array	
Using existing data drive in,.....	25
Using four drives in,	26

Using new drives in,.....	24
---------------------------	----

N

Networking.....	See Windows NT
E-mail notification	41, 45, 67
Firewalls.....	9
LAN/WAN	33, 34
SuperCheck requirements.....	9
TCP/IP	9
User administration	40

O

Object View	50
-------------------	----

P

Passwords	
Accessing Message Agent	65
Changing User,	75
Setting	76
Pull-Down Menus	
Connection	52
Help	53
Preference	53
View	52

R

RAID	See Arrays
Array Types	
RAID 0.....	104
RAID 0+1	105
RAID 1.....	105
RAID 3.....	105
RAID 5.....	106
Spanning	106
Management.....	107
Maximum capacity.....	See JBOD
Performance desktops	102
Safeguards	108
Rebuild	
Cancel	93
Manual.....	90
Using Rebuild Wizard	91
S	
Software Installation.....	9
Drivers.....	11, 12
NT services.....	33

Quick install	7	Halting	94
SuperCheck.....	33	Scheduling	69, 73
TCP/IP	9	T	
Spanning.....	27	Toolbar icons	56
Assigning drives to,	27	Tree View	
Status Bar	51	Collapsing	48
Stripe Block Size..... See Block Size		Displaying	48
SuperBuild..... See BIOS		Expanding	48
SuperCheck.....	31	Hiding	48
Adding users	40, 43, 65, 74	Refreshing display	49
Array info.....	89	Renaming levels.....	48
Array Synchronization	93	Using	48
Component installation	35	Troubleshooting	41, 109
Creating arrays	44, 79	U	
Deleting arrays.....	45, 89	Users	
Deleting users.....	40, 44, 74	Assigning Rights.....	75
E-mail notification.....	41, 45, 67	Changing Passwords	75
Enclosure monitoring	87	Creating	74
First use of,	40	Deleting.....	74
Installation scenarios	33	W	
Message Agent	44	Windows NT	3
Message Server.....	43, 60	Application server.....	103
Pull-Down Menus	51	File server.....	102
Setting password	76	Services	
System access.....	33	Message Agent	2, 32, 44
Toolbar.....	56	Message Server.....	2, 32, 43
User administration	43, 74	TCP/IP	9
User Administration.....	65		
Using SuperCheck	47		
Synchronization	107		
Arrays.....	93		

P/N: C6101ST66000000

NOTES